THE EVOLUTION OF THE CYBERSECURITY ECOSYSTEM

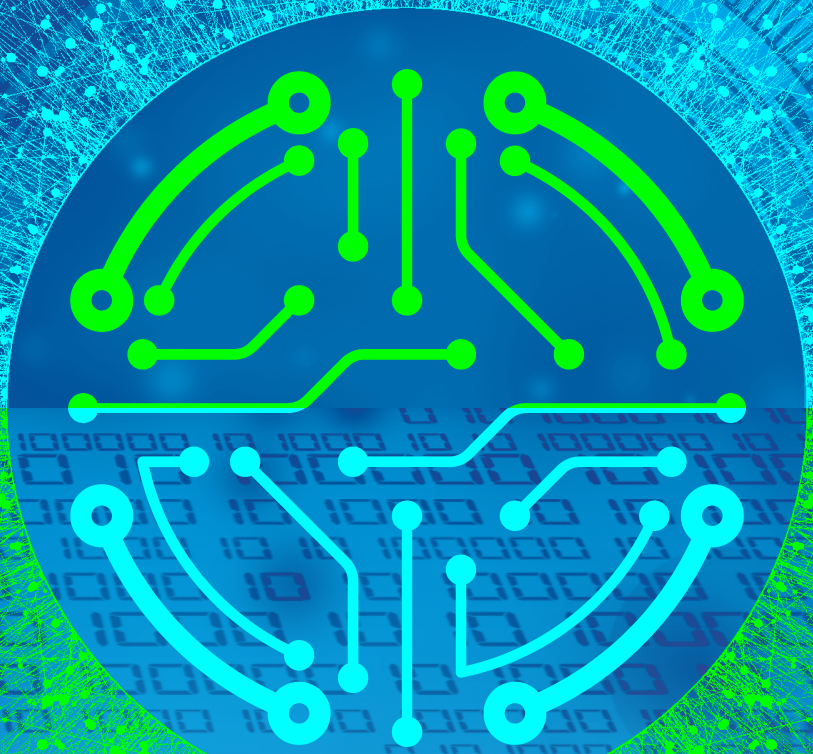

A CONVERSATIONAL JOURNAL FOR CYBERSECURITY AUTOMATION

CIS Center for Internet Security®

IACD COMMUNITY

LIST OF CONTRIBUTORS

Bret Bergman, COO Partners in Performance, America

Joe Brule, Chair OASIS OpenC2 Technical Committee

Geoff Hancock, CISO Advanced Cybersecurity Group

TK Keanini, Product Line CTO for Analytics Cisco

Philippe Langlois, Data Breach Investigations Report (DBIR) Author Verizon

Aubrey Merchant-Dest, Federal CTO Symantec

Kathleen M. Moriarty, Global Lead Security Architect EMC

Phil Reitinger, President and CEO Global Cyber Alliance

Shawn Riley, Chief Visionary Officer & Technical Advisor to the CEO at DarkLight

Tony Sager, Senior VP & Chief Evangelist Center for Internet Security

Kumar Saurabh, CEO LogicHub

Charles Schmidt, Group Leader The MITRE Corporation

Kathy Lee Simunich, Computer Scientist Argonne National Laboratory

Kim Watson, IACD Technical Director JHU/APL

Donnie Wendt, Security Engineer Mastercard

CISA
CYBER+INFRASTRUCTURE

We wanted to take a moment to thank NSA and CISA for all their support of IACD since its inception 5 years ago. All the work done on the IACD team, the growth of the IACD community, and this journal would have never been possible without the vision and support of our government sponsors.

TABLE OF CONTENTS

EVOLVING PERSPECTIVES

The State of the Operational Ecosystem

By Tony Sager and Phil Reitinger

To paraphrase our friend and colleague Rich Struse, Chief Strategist for Cyber Threat Intelligence at MITRE, on most days cyberdefense feels like "our people are chasing their robots"*.

In cyberspace, the bad guys have the upper hand: speed, anonymity, and leverage – essentially unbounded by space and time. They are also part of an integrated and automated criminal ecosystem featuring high return with low risk, easy access to attack toolsets, rental of infrastructure, global information sharing, and specialization (e.g., "money mules", reconnaissance data). The incremental cost to extend attacks is often minimal, while the likelihood of being caught is negligible. They disrupt our operations, steal our intellectual property, force us to spend vast amounts of money and manpower, and raise our uncertainty via a fog of botnets, criminality, and subverted Web sites.

Meanwhile the vast majority of our defenders are in effect pinned down by relatively mundane problems: poorly engineered software, missing patches, unenforceable policies, poor configuration choices, and inconsistent and sometimes conflicting security controls. They are asked to support ever-increasing demands for the business use of technology, and connection to partners with unknown security properties, at a pace that doesn't allow for thoughtful and secure integration. Technological developments like the Internet of Things increase the difficulty by orders of magnitude.

And yet there ought to be some "home court advantage" for defenders: control of their own IT; most attacks fall into a relatively small number of classes, taking advantage of a relatively small number of unique vulnerabilities over and over again; and fortunately (but sadly) there's enough experience in dealing with attacks to develop playbooks or reusable procedures. Moreover, with (still) more good devices than bad across the Internet and subsidiary networks, defenders have a powerful platform for observation of network and device behavior and distributed, automated response to malicious activity.

There seem to be two clear requirements. The first is widespread reduction of vulnerability, what we might call "secure by near default." Systems and software should be delivered and maintained in a secure state, and if the end user must take action, the effort required should be as close to zero as possible. Second, relying on our knowledge of common attacker tactics, techniques and procedures, and the availability of a broad sensing platform, networks should be able to defend themselves from the vast majority of threats, using automated collective defense.

---

* https://www.kuppingercole.com/watch/eic17_10_struse

All of this speaks to a need for much greater use of automation and standardization. And not just technology, but technology that is built directly into the architecture, made a natural part of acquisition, linked to policies, supported by training and operational processes, and adaptable to new information. We need all of this at a reasonable cost, built into commercial off-the-shelf products, and based on open industry standards.

For many of us, the roots of security automation go back to data, especially CVE and all the work that flowed from there into SCAP and related projects[†]. Abstractly, this meant standing back from the overwhelming cyber "Fog of More"[‡], and looking across the broad community of cyberdefenders to see the problems that every defender had to solve on their own. For example, in Tony Sager's part of the Defense Department, he watched DoD spend huge amounts of money and time just to collect, store, move, and reformat data from numerous sensors and other sources. By naming, numbering, and describing relevant things in open, machine-usable ways, we could start to build an environment of shared data, labor, and ideas. As the data that drives cyberdefense becomes more shareable and "frictionless", we can turn our attention to use of that data – building in prevention, response, playbooks, and training.

At its heart, cyberdefense is a decision-making, risk-managing machine, fueled by information. We need to move from managing information technology to managing information.

At its heart, cyberdefense is a decision-making, risk-managing machine, fueled by information. We need to move from managing information technology to managing information. To gain the defensive advantage, machines and people must be able to rapidly collect, correlate, and use information of many types and from many sources (e.g., IT components, network devices, specialty security tools, threat data) in order to assess the current risks to our operations and take both automated and particularlized action for prevention and response. These are all crucial defensive actions that need to be seen as part of a holistic cyberdefense machine that manages space and time to defensive advantage.

What's really clear – none of can solve this problem alone. Beyond the public-private bumper stickers, we need a community to emerge that includes security practitioners, researchers, buyers, operators, educators, IT and security vendors, and policymakers.

[†] https://nvd.nist.gov/

[‡] https://www.youtube.com/watch?v=OZLO-xekp3o

CONVERSATIONAL TOPIC: ECOSYSTEM

"A Shared Ecosystem"

By Aubrey Merchant-Dest

Cybersecurity and risk management require context and visibility that is relevant and timely, consumable and actionable, using a taxonomy that everyone understands. We want to communicate the 'so what and why should I care' and understand the 'what can I do,' with a measure of cost, potential risk and likely consequence. Cybersecurity and risk management are symbiotic and affect individuals and enterprises alike. Loss of sensitive data can manifest in ways only limited by imagination, and Artificial Intelligence (AI) opens yet another set of outcomes which we are only now starting to question and grasp. Ethics, integrity and provenance are where we should focus, this is the cyber defense and risk management challenge moving forward. We're approaching a point where loss of trust in the connected realm risks quality of life.

The majority of our focus in cyber defense thus far has been on identifying external threats, attempting to block them at the network perimeter or at some endpoint device. Data now resides in online applications, or shared services hosted externally. There's value or convenience for businesses and consumers respectively, but risk is never eliminated, only transferred and/or accepted. Recall that 'opportunity cost' is the measure of loss or gain from an alternative not chosen. This is the realm of the Actuarial for large businesses, but how do we resolve that when the data or object can be used to manipulate individuals, creating public harms? The Centers for Disease Control (CDC) may be the best example for how we identify and mitigate public harms relative to cyber security and risk management. They identify threats, communicate a course of action that is clear and concise, and define a protocol understandable by all entities. This is exactly the challenge faced in our networked, connected world.

People, devices, applications and data are the demand and supply-chain in today's global economy, and any one of these can introduce risks or threats to an environment. Challenges in risk and threat protection are exacerbated further as modern compute paradigms are adopted; for example, more distributed data requires more visibility and added controls/management. Automated approaches to code scanning, configuration management and threat protection are making a positive impact, though we still require correlated/converged visibility and contexts to get us to the next level. This could potentially be resolved through a ~universal taxonomy describing 'publicly shareable attributes' about risks and threats.

Person entities (PEs) and non-person entities (NPEs) are core to this equation, factored with anonymity and privacy concerns in mind. Consider medical devices and the data they collect and forward as a use case, what attributes are collected and correlated to mind your physical being? Are attackers targeting these or similar devices in another geography? That would be useful

information to share, elevating awareness and preparing for an outcome. We need to share contexts, not the proprietary information that creates or identifies mal-activity or vulnerabilities. We don't want to commoditize cybersecurity and risk loss of efficacy, but it must scale your existing workforce and blanket critical investments. Having a bigger picture view of everything relevant to your data property and/or community helps make for a safer environment.

If we took a CDC-like approach, we need to understand where threat is, what they are targeting and through which vulnerability, the type of endpoint and its hygiene, the targeted infrastructure, and other attributes useful for identifying the 'hot zones' or even pandemic cyber threats. In other words, is it viral, bacterial or fungal, and how is it contracted, what are the symptoms, and how is it treated or prevented?

In this approach, we don't need to define or standardize on all the data. We would just need to have a set of commonly understood pieces of information, and that information would have to be designed intentionally to allow people to easily map from the global issue to their personal situation.

This "edge computing" is a necessary component in the realm of "cyber physical systems" and therefore critical to risk management as we continue to innovate technical solutions. The new boundary is the data (and the services).

TODAY'S OPERATIONAL ENVIRONMENT

Organizations benefit from the agility and efficiency that today's environment promises, but still require the visibility and control possible with on-premises systems and tools. Organizational data will increasingly be housed within a Cloud Service Provider's (CSPs) data center, but not all risk is mitigated in their environment. The shared responsibility model requires entities to manage data "in" the CSP platform/shared service, which requires constant configuration management to mitigate both technical and human factors.

Compounding this is that the definition of an endpoint continues to evolve. The future of cloud computing will be far more distributed than we currently imagine. Smart cities, autonomous vehicles and the continued proliferation of connected "things" will force further distribution of compute fabric closer to endpoints/sensors to improve real-time analytics and take actions at machine speed. This "edge computing" is a necessary component in the realm of "cyber physical systems" and therefore critical to risk management as we continue to innovate technical solutions. The new boundary is the data (and the services).

CONTEXT IS EVERYTHING

There are many touchpoints supporting business continuity. Security or IT operations implement tools and technology that identify and mitigate risks which manifest from internal or external threats, but more context and visibility is needed to identify unusual behaviors which extend beyond the capability of traditional security stacks. Your data and applications extend to multiple properties, you need visibility and meaningful contexts to enable user and entity behavioral analytics (UEBA); to be clear we are referring to PE and

NPE correlation. Identity proofing all touch points is paramount, but having some well-defined, set facts about entities accessing your resources is critical to making decisions at machine speed.

While today's security tools can identify known and even novel threats, not all threats manifest as malware. Consider 'deep fakes' and the risk they present both personally and politically. We should consider asset management of digital information/objects similarly to entities (PE and NPE) that interact with it, and this is inclusive of communications from the network to the application layer protocols and the data itself. Applications are business oriented, networks provide connectivity, security tools detect and protect from threats, and they need to operate in unison, in cyber-relevant time. All of the interacting components can become sensors, essentially a detective at every corner to help assess, decide, and monitor. Sense-making of your operational environment becomes more achievable, and it can indeed span all of your data properties. We can now consider automating response and recovery procedures with confidence and reliability.

All of the interacting components can become sensors, essentially a detective at every corner to help assess, decide, and monitor.

EVERYTHING AND EVERYONE IS AN ENDPOINT

Humans, connected devices, hypervisors, containers and workloads are all endpoints. Applications and information determine the demand and supply-chain. That's a lot of touchpoints to support supply-side business operations, all necessitating consistent configuration management and policy control, and privacy/attribute protections. Security tools aid in identification and mitigation of risk or threats from internal or external sources, but context, visibility and analytics are necessary to track and reveal unusual behaviors which extend beyond the typical security stack. Today these two sets of information are not easily correlated to support more timely and effective decisions, there's a need for rapid convergence. Seemingly unrelated and even remote 'signals' when properly correlated can pinpoint high-value 'noise' early. Yes, we're talking SIGINT for a connected world, a CDC-like approach, with a threat intelligence overlay.

The real need is to log and detect activity in a manner that makes it easier to relate to the threat (in your context) to understand your exposure, risk, and mitigations, if you determine you are or may become infected. That requires a level of visibility and understanding that we currently do not have about our end-to-end system interactions and the transports connecting them. Networks, endpoints and cloud applications of today utilize APIs to generate log data that can assist to inform security operations, and they all provide sensor-like capability and can contribute to sense-making. The human factor may be the ultimate challenge, we are souls with built-in sensors, but how do we respond to the unknown and/or unusual? Think of a type of truth-table factored with experience or otherwise learned, reactive to stresses past or present. Think about the effect of a well-crafted phish or an urgent SMS from a loved one. Point is that a larger community can benefit from ~attribution

and relate it to a larger concern. Think about that in context of SIGINT, bigger picture intelligence. This approach should facilitate broader critical thinking about cyber awareness.

TODAY'S THREAT ENVIRONMENT

Irrespective of the size of your organization, threat protection and intelligence should be an integrated component of risk management and cybersecurity protection. It is useful to understand the threat landscape even when you don't have direct exposure. You gain insight on how your organization might be affected and the potential impact to systems health, and/or to understand how threats are evolving, where they are moving. You see and understand this early in the cycle and stay in front of a situation.

Mature security operations likely utilize both commercial and open source intelligence feeds for additional context and improved dimensionality. Public/private information sharing occurs with authorities when threats are egregious enough to warrant national security concerns or potential pandemic.

Put another way, the intelligence must be easily relatable and relevant, answering the so what and why should I care questions, the only questions actually that matter, and do so at a useful point in time.

Optimum information sharing requires a trusted-source, and trust can gain or wane. There's plenty of useful threat intelligence available if you're able to parse and correlate it to your environment. Put another way, the intelligence must be easily relatable and relevant, answering the so what and why should I care questions, the only questions actually that matter, and do so at a useful point in time.

BRIDGING THREAT AND NETWORK DEFENSE COMMUNITIES

As we've been leading up to, let's define what a CDC approach to cybersecurity and risk management would require. It is a protocol, actually, for risk management and intelligence collection and dissemination, sufficiently abstracted for wide comprehension. It then would support a ~universal plug-in architecture, a platform ideally, an integration fabric minimally; an open platform that gives you control over your enterprise telemetry and security data: how much you collect, how long you retain it, and where it resides. It also provides a standard, cross-product schema for analytics, reports, and dashboards. It links PEs and NPEs to applications and/or objects, everything contributes to sense-making.

It normalizes the event data that it collects from different sources to match a minimum, common standard, using the same attribute names and values for equivalent events. The schema is an extensible information model that defines one common taxonomy for the event data that is collected from different products. The schema is also a cross-product standard that is maintained and upgraded as necessary; it defines the event types and groups them in a number of categories.

Extensible – important to ensure interface points to ensure extension to know and understand capabilities, information, or policies. A starting point would be persons, devices, transports, applications, data, frequency, geography, etc.

Open, Orchestrated, Managed, and Protected Integration – as more and more sense- and decision-making are performed by analytics and algorithms, the orchestration needs to connect capabilities (not products and services, but maybe even packages of them to meet a capability) in a way that can be clearly articulated, synchronized, monitored, assessed, measured, and secured.

Maintain Trust – we must be careful to guard against possible negative outcomes as we depend more on automation and artificial intelligence. Trust is additive and subtractive, provenance must be well understood and controlled. Think multiple notary public roles, distributed and coordinated hierarchically.

We all play a role and have societal responsibilities, we need to adapt this to cybersecurity, realizing the potential for risk, contributing to early identification and risk mitigation for the greater good.

GIANT STEPS ARE WHAT WE TAKE (WALKING ON THE MOON)

Do you recall Bird Flu (Avian Flu)? The CDC website defines the basics, provides guidance and updates, as well as prevention, treatment and subtypes/variations. This answers the 'so what and why should I care' as well as the 'what can/should I do.' It does so in terms and contexts easily understood by the community at large, across the globe. Cybersecurity and risk management should be no different. The Department of Homeland Security implemented 'see something, say something' many years ago, and while you can argue its effectiveness, the intention is simple to understand. We all play a role and have societal responsibilities, we need to adapt this to cybersecurity, realizing the potential for risk, contributing to early identification and risk mitigation for the greater good.

Stop for a moment and consider the number of connected devices that support commerce and social communications, then consider the variables related to hygiene of devices and (un)intentions of human factors. It's difficult to wrap your head around the volume and variables. Technologists should believe technology is about humans and not gadgets. But bad actors see something very similar, opportunity to exploit technology to deliver an outcome, and they're playing us in ways that should offend us all.

Every connected device has some ability to report some activity, what is needed is a mechanism to correlate otherwise loose signals to identify activity of interest. This requires sense-making attributes to be converged and analyzed to include end-to-end communications. We can understand where threats are emanating and if they were stopped up or downstream. But that takes coordination across multiple entities (service providers, equipment vendors and others). We all have a view of cybersecurity and risk management, but it is generally constrained to specific areas of focus

or specialization. Network folks are concerned with uptime and availability, security operations have the very same concern but from a different vantage point. The cost is cheaper if we can block it at the network/transport layer. It's time to create an overlay model to gain additional visibility and contexts. Only then can we improve end-to-end situational awareness and improve our risk management game.

DELL EMC

Shared Ecosystem Push from Evolving Standards

By Kathleen M. Moriarty

Information security is in a transformational period. We have the opportunity to embrace change that could result in an overall improved security posture for both large and small organizations to improve the overall security posture of Internet connected systems. New threat models and advanced attacker techniques point to the need for a shared ecosystem model with holistic automated control management. Let's dive into the standards evolution that is unfolding to counter these threats and providing opportunity for change and advancement.

TRANSPORT SECURITY

Over the last twenty years, standards have evolved to focus on improving transport encryption security, with a recent (since 2013) increased focus on options for strong transport encryption that is easily deployed. Transport Layer Security version 1.3 (TLSv1.3) and related protocols like QUIC that leverage TLSv1.3 are provably secure (under reasonable assumptions), provide forward secrecy, and not only reduce the exposed meta data on the wire, but also prevent passive monitoring that was possible in earlier versions of TLS through the use of static keys. The recent efforts to automate certificate management (Let's Encrypt) via the Internet Engineering Task Force's (IETF) Automated Certificate Management Environment (ACME) protocol has eased deployment and management of certificates and can be tied to an increase in deployed web (HTTPS) encryption from about 30% to over 80% (https://letsencrypt.org/stats/).

Ultimately, the security gains of TLSv1.3 will improve the overall security posture; however, for organizations using primarily passive monitoring techniques to reduce threats on their networks or to perform troubleshooting, this trend towards stronger encryption that can't be passively intercepted is disruptive. To realize gain in security posture, alternate controls must be established to detect and prevent threats. Additionally, threat actor techniques are evolving. Threat actors are working hard to ensure that indicators of compromise used in one attack do not match the attack vectors or indicators in the next attack by recompiling code and altering techniques (https://www.forbes.com/sites/samcurry/2019/06/27/indicators-of-behavior-the-new-telemetry-to-find-advanced-cyber-attackers/#10668218193e). While in the past interception techniques have been somewhat effective, secure coding practices to prevent exploitation and security control monitoring for system and network changes may be more effective at countering these new threats longer term. Establishing a monitored baseline of controls that are actively monitored may be the best

way to detect threats early in the kill chain (https://www.lockheedmartin.com/en-us/capabilities/cyber/cyber-kill-chain.html).

ENDPOINT SECURITY AND AUTOMATED CONTROL ASSESSMENTS

As we know, attackers will find the easiest point of entry and then use both privilege escalation as well as lateral movement techniques to attain their desired goal in an attack. Maneuvering also enables the adversary to change goals, while establishing persistence, to take advantage of previously unknown targets of opportunity. With secure on-the-wire encryption being ubiquitously deployed, endpoints (https://datatracker.ietf.org/meeting/105/materials/minutes-105-saag-00)—including applications, devices, and systems—are the easiest targets. Standards have existed for some time to aid in the management and monitoring of security controls at endpoints, but with the exception of network based controls, they have not been universally deployed or easy to manage without the purchase of additional point products.

Cost and ease of deployment are critical considerations to achieving a shared ecosystem model where small and large organizations can practically deploy secure systems that are cost effective to manage and monitor.

The Security Content Automation Protocols (SCAP) from NIST and the endpoint work from Trusted Network Connect (TNC) were good starting points in the previous ecosystem where add on security was accepted as the norm. SCAPv2.0 (https://csrc.nist.gov/Projects/Security-Content-Automation-Protocol/SCAP-2-0) recognizes that deficit and is working to shift the model to embedded security at the endpoint, from the OS or application vendor that is measurable (automated control monitoring). The automated control assessments are through universal controls accessed via standards based protocols to a shared repository. Any evaluation system may access the repository control assessment results information for reporting and detecting changes in the network if authorized.

Automated control management is essential towards easing the costs of managing the security of an organization against a holistic security control framework. Cost and ease of deployment are critical considerations to achieving a shared ecosystem model where small and large organizations can practically deploy secure systems that are cost effective to manage and monitor. Scale is essential in this new ecosystem model. By setting security control policies with expected results, an organization is better able to detect unexpected changes and thus have early warnings of possible infiltrations or simply infractions to established policies and procedures.

The standards evolution is driving the need to view security control management as a shared ecosystem with responsibility shifting to points of control that scale better not only for management, but for organizations of all sizes. Security control management tied to a security control framework (NIST 800-53, ISO27001, or top X controls such as those from CIS) are difficult to attain when resources are limited for information security. Automation, originating from the supporting vendor, is necessary to overcome this hurdle.

Automation is best supported when it is built into the system, indicating that we should remove add on point products and in-line services in favor of built in security with the ability to easily monitor for changes.

In addition to work on security control automation, both at NIST in the SCAPv2 effort and the IETF's Security Automation and Control Management (SACM) (https://datatracker.ietf.org/wg/sacm/about/) working group (which overlap for several standards), there are additional standards development efforts that may aid in achieving higher levels of security while offering opportunities to reduce resources necessary for security management, called attestation.

ATTESTATION STANDARDS

Shifting the responsibility back to the creator also reduces the demand on security professionals as the work to secure applications and software rests with the producer as opposed to every consumer…

In attestation, simply stated, a module would be chained using digital signatures from modules required for its support, along with claims about the security posture that may be tied to a hardware root of trust. There is attestation work happening in several standards bodies, with an effort to coordinate work across these efforts. The IETF's Remote ATtestation procedureS (RATS) (https://www.oasis-open.org/committees/tc_home.php?wg_abbrev=openc2) is defining formats for signatures with claims using a shared architectural vision from the Trusted Computing Group (TCG) (https://tools.ietf.org/html/draft-fedorkow-rats-network-device-attestation-00). While this attestation work is further off, it could be very helpful in a move towards more secure systems in a shared ecosystem. How so?

Attestations come from the creator of software who is responsible for vetting their software prior to signing it and providing claims about that software. The software creators link their attestations to software in which their software or module rely and presumably have performed adequate testing, including for security vulnerabilities prior to providing the attestation. This shifts the responsibility back to the software creator for the security of their product, which is a better fit for the evolving ecosystem that cannot rely upon point products in the middle to detect and thwart threats. Shifting the responsibility back to the creator also reduces the demand on security professionals as the work to secure applications and software rests with the producer as opposed to every consumer being responsible in models that support the need for point products treated separately from the vulnerable systems and applications.

SECURITY CONTROL AUTOMATION AND THE CLOUD

Microarchitectures or software modules in serverless architectures will also greatly benefit from automated security control management to detect any variances from policy that may indicate exploitation of a security vulnerability for early detection. Chained attestations provide an assurance of trust that software has not been altered from a known state that will in the future fold into automated security control management baseline expectations. These should be expected security controls in serverless architectures of the not

too distant future. Any changes detected in software or monitored security controls that signal a deviance from policies should trigger an alert. The alert would ideally be communicated between various vendor independent products using an open standard such as OpenC2 (https://www.oasis-open.org/committees/tc_home.php?wg_abbrev=openc2) providing notice of an alteration from expected security controls or other anomaly. Automated communications on detected threats or variances from expected controls, can enable rapid investigation and thwart attacks early in the kill chain.

While secure transport encryption with strong authentication is changing architectures and supporting connections across networks (multi-cloud) and between modules or applications help achieve 'zero trust models', more is needed to support a shared ecosystem that is borderless. The full supply chain needs to be considered and this means manageable security that scales for organizations of all sizes. Chained attestations of secure software that provides assurance that software has not been altered should be one of many security controls monitored. Detecting or preventing compromises early through automated security control management with responsibility placed with the providing vendor may be one of the best methods to improve the overall security posture at scale going forward. Open standards are shaping a path forward, that if embraced, could result in an overall improved security posture requiring fewer resources to manage as a result in a shift of responsibility to inherently provide security.

OpenC2

By Joe Brule

The need for Coordination of Attack Response at Internet Scale is hardly a contentious topic; however, discussions, strategies and efforts towards these ends have limited utility should the approaches fail to consider the following postulate: 'Internet Scale' involves two attributes.

- Any information model must be **widely understood and unambiguous** (a semantic metric)
- A cyber response must occur within **cyber relevant time** (a temporal metric)

Engineering strategies, design principles and approaches must support or at least be consistent with this postulate if we are to achieve coordinated response at scale. We consider the following engineering principles to be consistent towards this goal.

- **Separation of Concerns**: Decouple the functional blocks within a cyber-defense system to the greatest extent practical.
- **Standards based Interfaces**: All inputs and outputs (i.e. the primitives) must be standards based.
- All designs and implementations are **public knowledge** (an extension of Kerckhoff's principle).

Cyber systems are subject to a global threat from adversaries that are increasingly dynamic and operate at machine speed. Modern cyber defense products tend to operate in isolation and often statically configured. The use of statically configured point defenses against a global attack surface is not tenable. Future systems need coordinated defenses operating in cyber relevant time.

Creating coordinated cyber defense systems in the absence of standards is impractical. The integration of a suite of monolithic products may result in redundant cyber defense functions, incompatible functions and capability gaps. The functional blocks within a given product may be tightly coupled with other functions and the may not be directly accessible by way of an API. Typically, integration efforts are expensive, require customized interfaces, and if tightly coupled, difficult to maintain or modernize.

The Open Command and Control (OpenC2) effort is a technical committee within the OASIS International Standards Body. The purpose of OpenC2 is to define a standardized language for command and control of cyber defense technologies and their first suite of specifications were released in July of 2019.

MITRE

Spotlight on Standards: Lessons Learned from SCAP

By Charles Schmidt

Today's adversaries operate at computer speed and vulnerable systems can be compromised within seconds. At these speeds, it is not feasible to depend upon manual processes to assess security posture or correlate findings between security tools. Instead, enterprise defenders need immediate awareness of changes to endpoint posture, and security tools need to support automated collation of their findings in order to provide defenders with the relevant context necessary to take action as soon as action is needed.

In response to this need, the Security Content Automation Protocol (SCAP) was first published by NIST in 2011, the result of multiple years of development and community collaboration amongst users, commercial vendors, and the government. SCAP sets out guidance for the coordinated use of several "component standards" to work together to support automated posture evaluation of enterprise endpoints. Today, SCAP and its component standards are used as part of many organizations' cybersecurity strategies.

However, despite some success, SCAP continues to face numerous challenges. The US government is a major user of SCAP, but SCAP has received less adoption among commercial companies. Many commercial security vendors either have not adopted SCAP in their tools or support it in ways that hamper interoperability with other vendors' tools. Production of SCAP content also remains a major challenge, with gaps in coverage and long lag times between product release and availability of assessment content. At the same time, while generally more timely than earlier practices, SCAP assessments remain periodic and don't provide the real-time insight into enterprise security posture that today's security administrators need. For these and other reasons, the current version of SCAP has fallen short of its goal to provide a common framework that provides a broad and dynamic collection of content, support for real-time and open data sharing between tools, and comprehensive coverage of enterprise security assessment needs.

As work begins on SCAP's first major revision, SCAP v2, it is important that all participants in this effort understand what has worked and what has not, so that SCAP v2 can leverage the best parts of SCAP v1 while addressing the issues that have held it back. We believe that a committed community of participants can help expand and enhance SCAP v2 to provide real-time assessment of enterprise security posture to improve detection and enable defenders to react quickly when alerted to adversary activities. You are invited to be a part of this effort by joining the SCAP v2 community and sharing your insights to create a better framework for automating enterprise security. For more information, visit https://csrc.nist.gov/Projects/Security-Content-Automation-Protocol-v2.

CONVERSATIONAL TOPIC: COMMUNITY

The Cyber Fed Model: Creating Communities of Trust – Lessons Learned

By Kathy Lee Simunich

Around the beginning of this millennium, as computer intrusions were gaining attention, an idea was conceived to share cyberthreat information from machine to machine. At Argonne National Laboratory (Argonne), around 2004, a team of cybersecurity specialists and software engineers designed and implemented an automated cyberthreat-sharing system called the Cyber Fed Model (CFM).* A fundamental design concept was to design *trust* and flexible distribution into the system from the start, in order to create flexible (and dynamic) communities for sharing cyberthreat indicators (CTI) between the various national labs and other U.S. Department of Energy (DOE) sites and plants. Even today, it is quite difficult to convince organizations to share their indicators of compromise, so creating trust communities was the first challenge to tackle. The DOE community consists of the labs, sites, and plants across the DOE complex. DOE is the Sector-Specific Agency for Energy, so another trust community consists of energy sector organizations, both public and private, as well as the DOE Power Marketing Administrations (PMAs).

CFM utilizes a set of web servers that physically separates the four trust communities of our clients, DOE, the energy sector, other U.S. government departments and agencies, and general (a catch-all for other private industry and non-government entities). Within each domain, client sites are grouped into "federations," and sites may be part of one or more federations. Federations can be defined hierarchically as well; sub-federations can represent different groups within an organization. In addition to being structured around organizations, federations can also be organized around shared interest areas (e.g., high-performance computers, oil and natural gas, advanced manufacturing). Non-browser-based client programs are set up at the member sites to automatically upload (publish) and download (subscribe) data.

To address the issue of trust, CFM has a flexible and dynamic distribution system, which gives a site control of its own uploaded information whenever it uploads data. For example, a DOE laboratory can tag its data to be shared with all the DOE sites in the DOE Federation, it might want to share only with other national laboratories, or it may want to share only to a specific site. The site may change these permissions for each individual upload, if desired. The CTI data is doubly encrypted in transit and remains encrypted at rest. Each site and each federation maintain a GPG (GNU Privacy Guard) key pair. The uploading site will encrypt its data using the public keys for the federations and/or sites with whom it wishes to share. The CFM domain server uses these keys to limit availability to only the designated recipients (i.e., members of

* The system was known as the "Federated Model for Cyber Security" until 2010 and was the recipient of the 2009 DOE Innovation award.

specified federations and individually specified sites), who will download the data on their next poll of the server.

Around 2013–2014, DOE (represented by Argonne) became a participant in the Enhance Shared Situational Awareness (ESSA) initiative, a multi-federal-agency consortium that includes federal cybersecurity centers. One task of the ESSA community was to create the Access Control Specification (ACS), which defines how data may be accessed and further shared once CTI information is shared outside an organization to a central repository. Full implementation of this specification will be needed before federal agencies will be able to fully share their CTI information. Another decision that resulted from the ESSA effort was to agree to use the emerging standard for formatting CTI, called STIX™ (Structured Threat Information eXpression), which is now an OASIS standard (Organization for the Advancement of Structured Information Standards).

Once an organization achieves the capability to share CTI, they face the challenge of metrics: how much data is being shared, with whom, what type, and how much of each type of data is being reported…

More recently, a new interagency and private-industry community is being created by the U.S. Department of Homeland Security (DHS) through the Cybersecurity Information Sharing Act (CISA) of 2015. Before CISA, Argonne, through the CFM, shared CTI with other federal agencies and private industry. At the inception of the DHS's Automated Information Sharing (AIS) system, CFM was one of the first participants to begin sharing CTI data.

Once an organization achieves the capability to share CTI, they face the challenge of metrics: how much data is being shared, with whom, what type, and how much of each type of data is being reported; as well as duplicate indicators vs. multiple sightings of an indicator. CFM stores the CTI data in a database as an encrypted file and, if it is shared with CFM, the individual indicators are extracted and stored as well. The file representation needs to be maintained so that the distribution can be correctly directed. The individual indicators are necessary to compile metrics on types of data and how much of each type exists, as well as any data analysis or searching that needs to be done for reporting purposes.

Upon upload, a file must go through a series of processes within the CFM server before final storage, and errors that result in quarantine or rejection may occur at various points. The first step is to verify that the upload is coming from a known and trusted source. Next, the system verifies that the upload contents are in a valid file format. Then the indicators are checked against a whitelist; they are individually removed if either whitelisted or in an invalid format. Last, in certain cases, the file and header metadata (e.g., who is the originator of the data) go through a source obfuscation process so that when the other recipients receive the data, they do not know where the data originated. This anonymization step was required before asset owners in the energy sector would sign up to become CFM member sites. It is also important when crossing certain relationship boundaries.

Another challenge for automated machine-to-machine sharing is knowing when something goes wrong, whether on the client side or the server side.

CFM uses a centralized logging system to log the status of the running systems. It logs when files are received, from whom, any metadata included in the headers from the uploading sites, and any warnings or errors in processing the files. Systems are set up to monitor the health of the various CFM processes and will email alerts on various conditions to the CFM team. They also monitor connections to external partners and send alerts on connection failures. Custom CFM monitoring processes continuously connect to each domain server and will email an alert if any machines are down or excessive errors occur.

Echo cancellation was also a difficult hurdle. For example, if a participating site both uploads data and downloads data, the uploader may not want the same data back that they previously uploaded. This becomes especially challenging if the server obfuscates the source of the data. For instance, Argonne had to add a marker in the STIX documents it uploads to the AIS feeds so that the marker could be checked when making the next download call, so as to not re-upload DOE data back to CFM. The CFM server implements a "what's new" algorithm that maintains a "bookmark" for every client, marking which files that client has already downloaded for each call. By default, this filters out the user's own data so that the user does not have to do it.

One of the lessons the CFM team learned was limiting our focus only on indicators. Even though it is the first step toward sharing, many indicators need to have some sort of context attached to them to be useful. . .

Interoperability became an issue once we started sharing outside the CFM platform. Argonne needed to create a custom-built translation layer that could not only convert from the XML-based format of CFM messages to the more industry-adopted STIX format, but also work with different authentication schemes. CFM utilizes Basic web authentication, custom validation processing, and GPG encryption, whereas DHS AIS uses PKI (Public Key Infrastructure) certificate-based authentication. Each platform may have a different authentication scheme. The CFM platform needed to support each of these in order to interact with the platform that uses them.

One of the lessons the CFM team learned was limiting our focus only on indicators. Even though it is the first step toward sharing, many indicators need to have some sort of context attached to them to be useful; users want to know what the quality of the data is, and what their organization's response to it should be. The STIX architecture was designed after CFM became operational, and one of its good features is that it tries to capture much more context that can be associated with an indicator, such as: course of action taken; tactics, techniques, and procedures (TTPs); threat actors; and campaigns. All of this is information that would help another analyst defend their organization's networks. Participants need to strongly adhere to standardization in representing these higher-level concepts, so that the meaning does not get lost in interpretation. Version 2.x of the STIX architecture has gone a long way toward rectifying the overly flexible STIX 1.x specifications, where there were multiple ways to encode the same CTI data. Version 2.x also captures relationships between observed indicators and their higher-level context elements. These relationships and the additional associated context will be critical to the evolving role of CTI sharing.

The future of CTI sharing needs to evolve by incorporating added context and relational data, sharing behaviors of adversaries, and aiding in the composition and sharing of "playbooks." Adding playbooks to CTI sharing will allow organizations to capture and orchestrate potential responses that can be automated to quickly and more thoroughly defend organizations. In addition, implementing a distributed search capability would move the focus of cyberthreat information sharing from a "publish" model to a "research" model. A distributed search system could allow participants to search unpublished data from an organization (such as full packet capture, or even cyber-physical data) so that an analyst or team can build up the context or playbook and then share that with the community. Analysts would be able to discover behaviors across multiple attacks and define proactive defenses.

Shifting to a behavior-based approach would focus on the bigger picture of the cyberthreat versus isolated pieces of the puzzle. Instead of focusing on a single action, analysts can start viewing attacks as a sequence of actions, possibly utilizing Mitre's ATT&CK matrix. Playbooks that define workflows for response or remediation against the behaviors/sequencing of TTPs related to individual campaigns and adversaries can be shared to protect the community at large.

Playbooks that define workflows for response or remediation against the behaviors/ sequencing of TTPs related to individual campaigns and adversaries can be shared to protect the community at large.

Solving these challenges will require collaboration and a coordinated effort. However, with a more comprehensive understanding of the threat, and the availability of automation and orchestration capabilities, analysts will be able to disrupt the adversary in ways that will cost significantly more time and effort to work around than today's typical response of blocking a single indicator. All of this will contribute to the long-term improvement of cyber-situational awareness.

JOHNS HOPKINS
APPLIED PHYSICS LABORATORY

Evolution of Cybersecurity Communities

By Kim Watson

I have been in cybersecurity and its predecessor fields for 30+ years, and started my foray into security automation sometime around the year 2000. Over this time period, I have noticed an evolution in the cybersecurity communities I have engaged with...

Looking back, I realize that these communities matured in stages. These stages, and the transitions between them occurred as: the field of cybersecurity matured, relationships between organizations became more interdependent, practitioner leadership roles shifted, and the community perspective broadened. This article provides insight into the stages of community and it highlights some strategies to get us to a place where ideas like collective defense may become a possibility.

STAGES OF CYBERSECURITY COMMUNITY

We participate in community because it offers us something that we value/desire. We are willing to meet the expectations and we trust that if we do our part then the community will deliver appropriate results.

The stages of community reflects a growth in expectations, returns, and required trust. Our role is more evolved and what we gain is more substantial, and we often become more dependent upon each other. Not all communities need to advance through all the stages to be effective. It really depends on the purpose of the community. As we mature our perspectives on cybersecurity and try to establish partnerships and complex relationships, we need to understand the underlying foundations of trust and expectations so we can establish and nurture an environment where the community can thrive at a more advanced stage.

STAGE 1: GUIDANCE

Purpose: To advance the state of cybersecurity.

Expectation: Participants were expected to contribute to the body of knowledge; to help define and implement best practices.

Return: In return, participants would have vetted and credible guidance to follow. The community guidance would provide evidence to management of value so they would support implementation.

Trust: The participant trusted that those in the community are credible.

STAGE 2: ADVOCACY

Purpose: To externalize and evangelize new or advanced perspectives on cybersecurity with stakeholders and governance bodies.

Expectation: Participants were expected to support the creation and adoption of norms for assessment and valuation of cybersecurity, in particular the best practices created by the previous community.

Return: In return participants would be able to make recommendations on investments and use of resources that would result in receiving credit for meeting norms.

Trust: The participant trusted that those developing and advocating for the standards shared their values/concerns. That they represented an agreed upon view of risk and consistent understanding of business needs/constraints.

STAGE 3: PARTICIPATION

Purpose: To share operationally-relevant information to jointly prioritize or characterize cybersecurity risks.

There is an understanding that the problem *is larger than any one entity. Participants . . . are willing to be accountable to expectations of the group that may be more restrictive or larger than what they are responsible for on their own.*

Expectation: Participants are expected to participate with integrity. To share what they know as openly and honestly as possible.

Return: In return, participants would have access to knowledge and capabilities that they could not develop alone.

Trust: The participant trusts that the others are providing unique, appropriate, and relevant information, methodologies and services.

STAGE 4: SHARED RESPONSIBILITY

Purpose: To consistently define and mitigate cybersecurity risks. There is an understanding that the *problem* is larger than any one entity.

Expectation: Participants are expected to participate in a manner that actually improves the value of what the community produces. They are willing to be accountable to expectations of the group that may be more restrictive or larger than what they are responsible for on their own.

Return: In return participants receive tailored knowledge that scales and is more consumable or embedded in the capabilities produced by the community.

Trust: The participant trusts that participation generates a greater benefit. That they would be worse off if they did not participate. That the result is greater because of the community investment and they could never match it on their own. The fact that they have transitioned through the other stages has built a foundation that makes it easier to extend this level of trust.

STAGE 5: PARTNERSHIP

Purpose: To develop and implement cybersecurity risk mitigations at scale. There is an understanding that the *solution* is larger than any one entity.

Expectation: Participants are expected to own particular parts of the solution space on behalf of others in the community.

Return: In return, participants are able to focus on a set of the solution space, and rely on others to provide services and perform functions on their behalf. This frees up local resources to be best in class in the area that is your responsibility.

Trust: The participant trusts that all the others will do their part, and do it well enough to meet the needs of the participant. This is more than the trust required for the previous stage, where it was just necessary to trust there was a benefit. At this stage, you have to trust that there will be no harm or negative consequences.

TRUE PARTNERSHIP

True partnership means you surrender something. There is something you are no longer doing that you have handed off to your partner and vice versa. That takes real trust.

True partnership means you surrender something. There is something you are no longer doing that you have handed off to your partner and vice versa. That takes real trust. Your bottom line, your success, your credibility, your reputation are in someone else's hands. That is rarely acceptable, but for a business or government entity it could be catastrophic if that trust is misplaced. So how do we get there?

It is possible to build the requisite trust, expectations, and capability/capacity by cultivating it in communities that are already in the shared responsibility stage. But this takes time, and I honestly don't believe this will happen until it is a necessity…until your bottom line, success, credibility, and reputation absolutely cannot exist outside of having others do things for you that you can no longer do for yourself.

Many of the ideas being discussed to protect and defend the nation involve partnerships. Unfortunately, the communities being asked to engage in these partnerships have not developed the necessary trust to meet expectations.

ADVANCING SHARED RESPONSIBILITY, CULTIVATING TRUST

Looking at the stages, and what it takes to transition, it seems that we need to find ways to make people feel that part of belonging to a community is to participate in a manner that improves the value of participation for all. The largest barrier to this type of contribution might be a fear of having misplaced their trust…that the effort and risk of the contribution will not result in the added value of the desired return. Worse yet, fear that if they trust inappropriately, that their very contribution could be used against them by stakeholders, regulators, auditors, etc.

If this is true, then employing Low-Regret strategies is a way to get community members to transition from a model of Participation to one of Shared Responsibility. Develop, employ, and share techniques to minimize the impact to the participating members, even if their trust is misplaced. Work

together to define ways to be successful even if community doesn't deliver fully on the expected results.

An existing example of this type of community, built on the establishment of Low-Regret techniques, is the open source community. Members take code developed by others and use it in their own environment. They employ test and evaluation techniques and openly share results of their analysis when problems or vulnerabilities are discovered. Organizations believe the open nature of the projects provides some validation of their trust in other members' contributions. But they also employ processes to minimize the effect on their business operations even if that trust is misplaced.

Over time, small groups that have successfully implemented a Shared Responsibility model may develop the trust required for a small number of Partnerships to form organically.

Over time, small groups that have successfully implemented a Shared Responsibility model may develop the trust required for a small number of Partnerships to form organically.

CONCLUSION

We can continue to mature our community to the Shared Responsibility model, particularly by identifying strategies and techniques that compensate for the required increase in trust by minimizing the operational impact associated with misplaced trust. Whether we will ever get to a true Partnership model will depend on external factors that influence perception and social expectations/accountability. Of course, if we were really operating from a model of Shared Responsibility, maybe the natural evolution of pockets of Partnership would be enough to support some implementation of collective defense.

CONVERSATIONAL TOPIC: BUSINESS

Crossing the Great Divide

By Bret Bergman

Cybersecurity highlights one of the biggest corporate communication challenges of our time – bridging the gap between business and technical leaders. I call this the great divide. As OT/IIoT becomes pervasive, cybersecurity ("cyber") is increasingly being recognized as the business risk it truly represents, especially for industrial companies who are just now fully arriving at the cyber party. With the proliferation of IIoT, predictive analytics and autonomous vehicles, the divide is becoming more acute and can materially impact the bottom line.

As someone who has spent more than 30 years in business and technology in a largely industrial context, I've seen the issue from all perspectives. The divide is driven by a variety of factors that can be grouped into two buckets:

- *Orientation/language*– the business side is focused on earnings and growth and talks in terms of digital transformation; whereas, the technology side focuses on assessing, installing and maintaining the IT infrastructure and talks in terms of cloud migration, firewalls and NIST audits.
- *Personality types*– it's overly simplistic, but it comes down to one side being comfortable with face-to-face verbal communication and the other preferring any other mode (email, texts, anything electronic).

As cyber pervades the day-to-day operations of the business, technical leaders find themselves in the limelight in a way they never were before. Once cyber becomes a business imperative, technical leaders regularly meet with the CEO and COO, people with whom they may have traditionally met once a year at most.

Here is a very simple, three-step process to begin your journey of closing the divide:

1. Establish a common ground– acknowledge the gap and establish that cyber is a business imperative and needs to be managed as such. Creating a high-level dashboard that both sides understand and in which all find value is a good step.
2. Leverage safety analogy– for industrial companies (to whom this is targeted) cyber has many parallels to safety, and leveraging those can be very helpful for all parties. A note of caution…use this analogy carefully as safety is a life-or-death issue, and cyber is not (except in very rare cases).
3. Utilize a facilitator– both sides can benefit from a respected "third party" that understands and is respected by both sides. This is often someone involved in transformation (from digital or traditional change management) who speaks both languages.

Good luck crossing that divide!

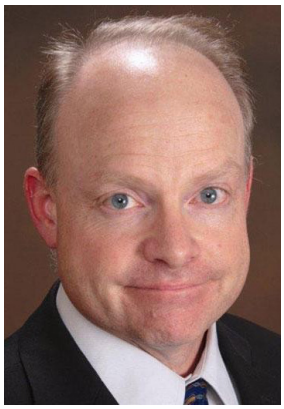

ADVANCED CYBERSECURITY GROUP

Building Business Resilience through Cyber Automation

By Geoff Hancock

Cyberattacks continue to increase in frequency and intensity; companies can't keep up. It's become increasingly evident that traditional methods, like anti-malware software, are no longer sufficient to keep sensitive data safe. And IT personnel are no match for the sheer volume of such intensive, sustained attacks. People's abilities to manually address such attacks and to make quick, accurate, highly impactful decisions are very limited. To deal with this situation, IT and business executives are finding ways to strengthen their cybersecurity strategy using automation. When automation is applied, repetitive, time-consuming actions can be performed more rapidly and in a more repeatable manner. This gives analysts time to focus on other, more valuable tasks.

Business resilience is an organization's ability to adapt to changes in business, good or bad while maintaining operations and protecting data, services, employees, customers, assets, and overall brand reputation. Enterprise risk management—which covers business, IT, and cybersecurity—gives the company a clear plan on how to do this. As part of that plan, the focus on cyber resilience involves *the ability to prepare for, respond to, and recover from a cyberattack* by tying together business risk with the capability to secure your digital assets. By reducing human intervention, cybersecurity automation also enables more time and energy to focus on essential elements of security operations. Some of the challenges security operations face today:

- **Security operations centers (SOC) are struggling with response times.** Organizations receive hundreds to thousands of threat alerts daily, and security analysts are unfortunately only able to investigate a small portion of these, making it difficult to manage and respond to real high risk, high impact problems.
- **Reduce the analysis of false data. There are many tools cyber professionals use to protect the organization. Many of them** create a lot of data. Data that is not prioritized or organized according to high or low-risk areas of the business. This tidal wave of data can overwhelm analysists and distract them from what is important.

Conventional methods are no longer enough. Security teams need automation because it provides fast and reliable detection of cyber threats. Still, implementing automation is not without its challenges:

- **Loss of control.** In many instances, the biggest obstacle to automation is simply a perceived loss of control. In reality, the right automation tool can provide a higher level of visibility and enhanced oversight of the entire cybersecurity process.

- **Lack of trust.** It's easy for a highly skilled human to feel as though he is more capable of managing incident response than a machine. Distrust of technology can be a huge hurdle to overcome, but ultimately—given the shift in type, frequency, and complexity of cyberattacks—it's a futile argument.
- **Fear of change.** One misconception is that automation spells the inevitable elimination of the human workforce. Will technology take over cybersecurity operations? Will robots replace the IT department? The fact is that while automation is undoubtedly changing the way people work, it also creates new opportunity for people to coordinate and use technology more effectively.

ADVANTAGES OF USING AUTOMATION

The fact is that while automation is undoubtedly changing the way people work, it also creates new opportunity for people to coordinate and use technology more effectively.

Cyber automation enables organizations to find and reduce risk in areas that may have a low but essential impact on the company. Streamlining these important risk management tasks frees up IT resources and staff to assess high-risk areas and helps them manage risk more efficiently.

- **Better decision making**. One challenge corporate leaders face is having to make critical business decisions quickly, often without being able to assess the risk to other parts of the business. Using cyber automation across the company provides a grouping of high, medium, and low-risk areas, each of which requires a direct correlation to business impact. This enables executives to make more informed decisions about how the business is run and where to make investments.
- **Increased efficiency**. Adding automation into IT systems management can help streamline business workflows, data management, and protection to create a much more stable and efficient environment.
- **A clear view of enterprise risk management**. Several of the high profile breaches in 2018 and 2019 have occurred because of poor enterprise risk management and the inability to identify high risk and low-risk areas. Not patching an IT system promptly may seem trivial on the surface, but if not done, can impact the company adversely. Equifax is a prime example of this. And not creating a plan that coordinates older IT systems with newer technologies to provide innovative business solutions can also damage a company, as in the case of Marriott hotels.
- **Focus on high-value activities**. The time saved through automation will free up cyber teams so they can allocate their energies to monitoring, threat mitigation, and response on the most critical areas of the system.

- **Minimize the risk of human error**. Automation can also help reduce the possibility of injecting human error into security tasks. For example, office networks are commonly compromised through phishing emails, which try to trick recipients into clicking links to malware. Phishing emails are becoming more complex, making normal analysis more challenging. Automated tools could quickly weed out such emails from company servers.

CONCLUSION

Today, low-risk tasks can be automated, which frees up resources to focus on high-risk areas that need extra protection and management. But, in the future, as cyber defense automation technology matures, business and IT risk management will benefit even more when higher risk areas become automatable. In general, cybersecurity automation can strengthen a company's risk management and increase resilience in the following ways:

. . . in the future, as cyber defense automation technology matures, business and IT risk management will benefit even more when higher risk areas become automatable.

- Methodically and persistently identifies the risks surrounding your business activities.
- Assesses the likelihood of an event occurring.
- Provides an understanding of how to respond to these events.
- Puts systems in place to deal with the consequences.
- Monitors the effectiveness of your risk management approaches and controls.
- Improves decision-making, planning, and prioritization.
- Helps manage capital investments and resources more efficiently.
- Helps you identify the highest risk areas and helps focus on ensuring those systems are resilient.

verizon

A Case for Standardizing Tooling Capabilities Language

By Philippe Langlois

For the defenders supporting and protecting networks, the deluge of data, alerts, best practices, notices and regulations can overwhelm even the hardiest of us. Fortunately for the defenders the tools they use have evolved in maturity, functionality and interoperability, however, understanding the specific capability of the tools and how they help you achieve certain regulations or best practices still presents a challenge. Does your vulnerability management platform also count as your organization's asset inventory? Does your Active Asset Inventory tool also help you create your list of active ports and services? Does your implementation of open source tools meet the same functional capabilities of an enterprise grade tool? What additional value does [insert new marketing buzzword] provide to protecting your organization?

One way that we can attempt to address these types of questions is through the development of a standardized language to describe cybersecurity tooling capability. We have standardized languages for describing software vulnerability (CWE, CVE), standardized languages to describe platforms and software (CPE) and also some growing standardization for understanding attacker methodology and campaigns (STIX, ATT&CK). While the task may seem insurmountable, there are so many tools that address problems in unique different ways. However, if the focus is not on HOW the tool functions, but on WHAT the tool does, the scope starts to narrow down drastically. With this approach we could start to whittle down the large corpus of tooling descriptions and distill them into more manageable classes of capabilities.

ONE POSSIBLE OPTION

As part of that distillation, the key point would be finding out the "WHAT", which is the action or activity that the tool is implementing/completing, and a good starting point is using the CIS Controls to describe the specific capability that is required for the sub-control to be implemented. While going through the process a pattern was discovered, many of the sub-controls had a relatively consistent structure, with one of a handful (approximately 12) Actions being applied to different subjects. Certainly not an extensive list or a perfect one, but intended to be a starting point.

- **Parse**: read and interpret different files and data into a consistent format
- **Scan**: Examine parts of something to detect specific feature
- **Store**: Write values and data into a queriable format
- **Block**: Prevent a specific event from occurring

- **Identify**: Uniquely represent data points based on a specific premise or criteria
- **Authenticate**: Verify something's identity based on a specific criterion
- **Patch**: To mend a weakness or vulnerability
- **Rate**: assign a standard or value to something according to a particular scale
- **Verify**: Demonstrate the accuracy of assumptions
- **Log**: make a recording of specific events
- **Set**: Turn on a setting or add a value to the system
- **Encrypt**: convert information into a cipher to prevent unauthorized access

Below is an example of how these action terms can be used:

- **Scan** the network for live assets,
- **Scan** assets for deviations to baseline,
- **Scan** assets for known malicious file hashes,
- **Block** from the network assets w/o legitimate certificates
- **Block** from executing unapproved applications
- **Block** from connecting unapproved connections.

The relationships would be key to understanding how specific capabilities tie back to different elements of the environment and how those protections roll up (or don't) in the environment.

In these extremely simple examples, we're looking at capabilities being described as the relationship between an Action and the Subject. The subject allows us to understand where the action is being applied and within which context. For example are we scanning the network for live assets, or are we scanning for known outdated software? The problem then becomes, what ultimately describes the subject. This is where there is a significant gap in my opinion. Ideally a standardized language of system subjects would allow for the capturing of relationships between subjects (a Filesystem is on a system, a network consists of different systems, applications can be distributed across different systems). The relationships would be key to understanding how specific capabilities tie back to different elements of the environment and how those protections roll up (or don't) in the environment.

However, in the imperfect world of technology I believe we need to introduce one additional piece, the "Scope". As much as we'd like to say that our capability would be applied universally across our environment it may only be available to certain platforms, to systems that are domain joined, or to systems that have agents installed. This scope would then help you understand the overall coverage of your organization's security tools. For the more mathematically inclined here is a quasi-mathematical way of looking at it:

Capabilities = f(Action, Subject, Scope)

By having a (this is but an option) standardized approach to describing capabilities we could theoretically make direct connections between security Tooling, security requirements and risks. This is obviously a very simple model

and one that needs significant work and expertise to advance from napkin scribbles to something robust. Hopefully we can start collectively shining a little extra light into the "fog of more" and get a better sense as to what are the capabilities that we're looking to purchase, that we need to purchase and that we're not currently take advantage of.

CONVERSATIONAL TOPIC: OPERATIONS

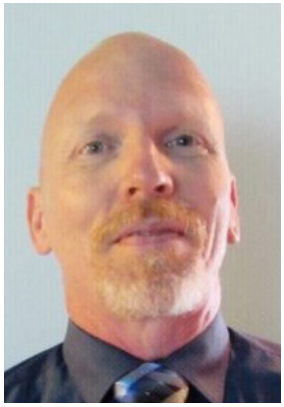

Driving Forces for Security Automation

By Donnie Wendt

Today's cyber defenders find themselves at a disadvantage despite technological advances in cyber defense. Among the chief causes of this disadvantage is the asymmetry in a cyber conflict that favors the attacker [1,2,3]. The increasing sophistication of the attacks increases the defenders' disadvantage. Finally, organizations face a growing shortage of cybersecurity professionals to meet the increasing demand [4,5].

ASYMMETRY AND THE ATTACKER'S ADVANTAGE

Cyber attackers have the advantage because the attackers only need to exploit a single vulnerability whereas the defender has the much costlier task of mitigating all vulnerabilities [1,2,3]. Attackers can choose the time and place of the attack which further disadvantages the defenders [6]. The ease with which an attacker can acquire and use an exploit coupled with the low likelihood of detection favors the attackers [7]. Once inside a network, individual actors in the cyber domain can have an asymmetric advantage and possess highly dangerous capabilities [8].

The typical use of homogenous platforms for information systems by many organizations can significantly increase risk. The use of similar operating systems, hardware, and applications increases the reward for attackers who can develop exploits that target the vulnerabilities in dominant systems [6]. This static nature of systems and defenses contributes to the imbalance that favors the attackers. An attack that exploits a vulnerability in a popular software application can infect millions of machines [2]. Attackers can install and analyze local copies of available cyber-defense applications and tools to discover the weaknesses and how to avoid detection.

Due to adaptive threats and rapidly changing technology, organizations make decisions about cybersecurity investments with imperfect and incomplete knowledge. Instinct, experience, and informed judgment are necessary for the prevention of, detection of, and response to cyber threats [9]. However, companies often must navigate lengthy, bureaucratic processes to implement new security technology [3]. Attackers can implement, analyze, and use new technology immediately.

Well-known, static defenses are increasingly vulnerable to threats from well-resourced attackers engaged in targeted attacks [10]. The predominately static nature of cyber defenses often requires time-consuming processes to reconfigure if they can be reconfigured at all [11]. The time required to reconfigure security devices in response to an attack allows the attacker time to locate and exploit vulnerabilities. The study of adaptive cyber defenses seeks to address this asymmetric advantage enjoyed by the attacker.

THE INCREASING SOPHISTICATION OF ATTACKS

A defender must first detect an attack before a response is possible. The increasing sophistication of attacks makes the identification of both successful and unsuccessful attacks more difficult. The detection of the attack should occur as early in the cyber-attack lifecycle, or cyber kill chain, as possible to minimize the ramifications of the attack [12]. Many of the sophisticated attacks, once inside a compromised network, seek to remain persistent. Such attacks are referred to as advanced persistent threats (APT). With an APT, the initial attack attempts to establish persistence from which to operate and call out to a command-and-control system [13]. The attacker can establish this persistence because organizations are often not aware of what software the organization has installed and running on each device. Even a device such as a printer can serve as the initial beachhead from which an APT can operate.

These technological advances allow attackers to not only develop more advanced attacks, but also to decrease cost, time, and risks associated with launching an attack.

The financial industry is a leading target of APT threat actors who intend to steal high-value data [3]. Attackers who invest in an APT are highly motivated and will devote significant time to compromise a target to achieve a specific goal. Advanced persistent threat actors will map out multiple paths to reach the target and pivot their attack as necessary to reach the end goal [13]. With the expanding complexity of systems, organizations present an increasingly large attack surface. The greater the perimeter, or attack surface, the more opportunities for the attacker to penetrate the perimeter and establish a persistence within the environment [13]. Current signature and anomaly-based detection tools have not been fully successful in detecting APTs. Detection of APTs by either signature or anomaly detection methods is challenging because attackers craft APTs for a specific target and often use unique attack vectors [14].

In addition to the increased sophistication of attacks, the tools and techniques used by attackers are more advanced. The increased use of automation on the attacking side, including management platforms and autonomous botnets and viruses, increases the difficulty for traditional defenses to detect and mitigate the attacks [15]. These technological advances allow attackers to not only develop more advanced attacks, but also to decrease cost, time, and risks associated with launching an attack.

NEED FOR SECURITY AT CYBER SPEED

Current human-centered cyber defense practices cannot keep pace with the speed and pace of the threats targeting organizations [16]. Further, the speed of attack versus speed of response gap is getting worse [17]. Defenders need to drastically increase the speed of both detection of and response to cyber-attacks. Organizations must automate many risk-based decisions to facilitate this increase in detection and response speed [17]. The human involvement must become more oversight and less direct involvement in the detection and response. The role of humans must shift from being predominately *in-the-*

loop to being *on-the-loop*. With this shift, humans will review and validate conclusions based on machine-learning and artificial intelligence [18]. Increasing the speed and efficiency of detection and response also requires rapid exchange of threat and incident detail among the automated defense systems. Such rapid exchange will require interoperability between systems at the technical, semantic, and policy levels [17].

CONTINUING DATA BREACHES

The number of recent cyber-attacks and the media attention given to those attacks gives the impression that such attacks are increasing in frequency, becoming more organized, and are more damaging [19]. Many advanced and well-orchestrated cyber-attacks have targeted industry, military, and government infrastructures with the main goal being the exfiltration of data [14]. An example data breach involving a US company resulted in the theft of 40 million credit card numbers and associated personal information [20]. The direct costs associated with the damages and recovery from the breach totaled $61 million. The breached company also experienced a 46% drop in profit in one quarter.

Perhaps the chief driver of security automation derives from the shortage of cybersecurity professionals to deal with the increasing threats.

Studies related to the costs of data breaches often look to quantify the direct costs of data breaches. If one considers the indirect costs, such as decreased profits and sales reductions, the actual costs of a data breach are likely much higher [20]. The average cost of a data breach continues to rise due to the increased frequency of cyber-attacks, increased remediation costs, and increased detection costs [21]. Large-scale breaches of data within the financial industry involving APTs are likely to continue. Cyber-attacks will remain a significant problem for financial institutions due largely to the complexity of the Internet and connected systems [22].

THE SCARCITY OF CYBERSECURITY PROFESSIONALS

Perhaps the chief driver of security automation derives from the shortage of cybersecurity professionals to deal with the increasing threats. The shortage of people with the requisite cybersecurity knowledge, skills, and abilities threatens to undermine the security of the systems upon which the financial industry relies and erode consumer confidence and trust in the financial institutions [23,24]. In a survey of security leaders by the Center for Strategic and International Studies (CSIS), 71 percent of the respondents stated that the cybersecurity skills shortage causes direct, measurable damage [25]. The same survey found that 25 percent of the respondents claim to have lost proprietary data through a cyber attack due to the cybersecurity skills gap.

Many sources report on the cybersecurity skills gap. The National Initiative of Cybersecurity Education (NICE), a program of the National Institute of Standards and Technology (NIST), found that there were 350,000 cybersecurity job openings in 2017 in the United States alone [5]. The

shortage of cybersecurity professionals shows no signs of improvement in the near term. A 2015 study by International Information Systems Security Certification Consortium, better known as (ISC)², predicted a shortfall of 1.5 million cybersecurity professionals by 2019 [4]. Additional studies have ranged from a predicted current shortfall of one million cybersecurity professionals [26] to a shortfall approaching 3.5 million by 2021 [5]. The cybersecurity profession is not keeping up with the increased demand [24].

The increasing sophistication of cyber-attacks capable of avoiding detection and the increasing frequency of cyber-attacks are reasons for the continued increase in demands for cybersecurity professionals. Another critical reason for the cybersecurity demand is the ever-increasing information technology (IT) footprint [4]. The expansion into mobile devices and cloud environments in conjunction with an increasing array of security technologies are major drivers for the IT expansion. The need to secure an expanding perimeter with more security tools spreads already scarce cybersecurity resources even thinner.

In addition to an increase in security professionals, the cybersecurity skills gap requires proactive threat hunting facilitated by advanced analytics, real-time threat awareness provided by comprehensive intelligence, and security architectures that are integrated.

A sign of the increasing scarcity of security professionals is increasing salaries [4]. The shortage of cybersecurity talent has led to increased compensation for cybersecurity professionals. Within surveyed countries, the median salary for cybersecurity jobs is at least 2.7 times the average wage [25]. In the US, cybersecurity jobs pay an average of nine percent more than other IT jobs.

Rising employee churn can also signal an increasing shortage of security professionals. The cyber workforce may be facing a burnout factor resulting in employment churn [4], and security operations centers (SOCs) are perhaps the hardest hit by burnout and employment churn [27]. Security analysts working in SOCs have unique skills and must operate in high-pressure situations to quickly analyze security events, decide on the response, and act to protect the company. Security analysts in the financial services industry face constant cyberattacks putting them under constant pressure to perform.

The cybersecurity skills gap likely cannot be addressed simply by adding more cybersecurity professionals. In addition to an increase in security professionals, the cybersecurity skills gap requires proactive threat hunting facilitated by advanced analytics, real-time threat awareness provided by comprehensive intelligence, and security architectures that are integrated. Though people with untapped cybersecurity potential do exist, the number of people capable of performing in a cybersecurity position effectively over time is likely limited [23]. Even if all viable candidates entered cybersecurity there might still be a significant shortage unless the demand for cybersecurity professionals can be contained. Technological advances in security and the use of automation can help address the demand side of the equation.

REFERENCES

[1] Carter, K. M., Okhravi, H., & Riordan, J. (2014). Quantitative analysis of active cyber defenses based on temporal platform diversity. *OALib Journal.* Retrieved from http://arxiv.org/abs/1401.8255v1

[2] Okhravi, H., Streilein, W. W., & Bauer, K. S. (2016). Moving target techniques: Leveraging uncertainty for cyber defense. *Lincoln Laboratory Journal, 22*(1), 100-109. Retrieved from https://pdfs.semanticscholar.org/15ea/51017d7395fd9cddd626704d1fc82fc42e3e.pdf

[3] Tounsi, W., & Rais, H. (2018). A survey on technical threat intelligence in the age of sophisticated cyber attacks. *Computers & Security, 72*, 212-233. doi:10.1016/j.cose.2017.09.001

[4] Suby, M., & Dickson, F. (2015). *The 2015 (ISC)2 Global Information Security Workforce Study.* Mountain View, CA: Frost & Sullivan. Retrieved from https://www.boozallen.com/content/dam/boozallen/documents/Viewpoints/2015/04/frostsullivan-ISC2-global-information-security-workforce-2015.pdf

[5] Morgan, S. (2017). *Cybersecurity Jobs Report: 2017 Edition.* Herjavec Group. Retrieved from https://www.herjavecgroup.com/wp-content/uploads/2017/06/HG-and-CV-The-Cybersecurity-Jobs-Report-2017.pdf

[6] Winterrose, M. L., Carter, K. M., Wagner, N., & Streilien, W. W. (2014). Adaptive attacker strategy development against moving target cyber defenses. *ModSim World* (pp. 1-11). Hampton, VA: ModSim World.

[7] Zheng, D. E., & Lewis, J. A. (2015). *Cyber Threat Information Sharing: Recommendations for Congress and the Administration.* Washington, DC: Center for Strategic & International Studies. Retrieved from https://www.csis.org/analysis/cyber-threat-information-sharing

[8] Rivera, J., & Hare, F. (2014). The deployment of attribution agnostic cyberdefense constructs and internally based cyberthreat countermeasures. *6th International Conference on Cyber Conflict* (pp. 99-116). Tallinn, Estonia: NATO CCD COE Publications. doi:10.1109/CYCON.2014.6916398

[9] Garvey, P. R., & Patel, S. H. (2014). Analytical frameworks to assess the effectiveness and economic-returns of cybersecurity investments. *IEEE Military Communications Conference* (pp. 136-145). Baltimore, MD: IEEE. doi:10.1109/MILCOM.2014.29

[10] Soule, N., Simidchieva, B., Yaman, F., Loyall, J., Atighetchi, M., Carvalho, M., . . . Myers, D. F. (2015). Quantifying & Minimizing attack surfaces containing moving target defenses. *Resilience Week.* Philadelphia, PA: IEEE. doi:10.1109/RWEEK.2015.7287449

[11] Zhu, M., Hu, Z., & Liu, P. (2014). Reinforcement learning algorithms for adaptive cyber defense against Heartbleed. *Moving Target Defense* (pp. 51-58). Scottsdale, AZ: ACM. doi:10.1145/2663474.2663481

[12] Fonash, P. M. (2012). Identifying cyber ecosystem security capabilities. *CrossTalk* (September/October), 15-22. Retrieved from https://secwww.jhuapl.edu/IACD/Resources/Reference_Materials/Resilient_Cyber_Ecosystem_Capabilities.pdf

[13] Byrne, D. J. (2015). Cyber-attack methods, why they work on us, and what to do. *AIAA SPACE 2015 Conference and Exposition* (pp. 1-10). Pasadena, CA: American Institute of Aeronautics and Astronautics. doi:doi.org/10.2514/6.2015-4576

[14] Virvilis, N., Serrano, O. S., & Vanautgaerden, B. (2014). Changing the game: The art of deceiving sophisticated attackers. *6th International Conference on Cyber Conflict* (pp. 87-97). Tallinn, Estonia: NATO CCD COE Publications. doi:10.1109/CYCON.2014.6916397

[15] Atighetchi, M., Benyo, B., Eskridge, T. c., & Last, D. (2016). A decision engine for configuration of proactive defenses: Challenges and concepts. *Resilience Week* (pp. 8-12). Chicago, IL: IEEE. doi:10.1109/RWEEK.2016.7573299

[16] Johns Hopkins Applied Physics Laboratory. (2016). *Integrated Adaptive Cyber Defense (IACD) Baseline Reference Architecture.* Laurel, MD: Johns Hopkins Applied Physics Laboratory. Retrieved from https://secwww.jhuapl.edu/IACD/Resources/Architecture/IACD Baseline Reference Architecture - Final 0PR.pdf

[17] Fonash, P., & Schneck, P. (2015, January). Cybersecurity: From months to milliseconds. *Computer*, 42-50. doi:10.1109/MC.2015.11

[18] Willett, K. D. (2015). Integrated adaptive cyberspace defense: Secure orchestration. *International Command and Control Research Technology Symposium.* Annapolis, MD. Retrieved from https://pdfs.semanticscholar.org/a228/81b8a046e7eab11acf647d530c2a3b03b762.pdf

[19] Cavelty, M. D. (2014). Breaking the cyber-security dilemma: Aligning security needs and removing vulnerabilities. *Science and Engineering Ethics, 20*(3), 701-715. doi:10.1007/s11948-014-9551-y

[20] Soomro, Z. A., Shah, M. H., & Ahmed, J. (2016). Information security management needs more holistic approach: A literature review. *International Journal of Information Management, 36*, 215-225. doi:10.1016/j.ijinfomgt.2015.11.009

[21] Joyce, A. L., Evans, N., Tanzman, E. A., & Israeli, D. (2016). International cyber incident repository system: Information sharing on a global scale. *International Conference on Cyber Conflict* (pp. 63-68). Washington, DC: IEEE. doi:10.1109/CYCONUS.2016.7836618

[22] Zheng, R., Lu, W., & Xu, S. (2015). Active cyber defense dynamics exhibiting rich phenomena. *HotSoS.* Urbana, IL: ACM. doi:10.1145/2746194.2746196

[23] Cobb, S. (2016). Mind this gap: Criminal hacking and the global cybersecurity skills shortage, a critical analysis. *Virus Bulletin Conference* (pp. 1-8). Denver, CO: Virus Bulletin. Retrieved from https://www.virusbulletin.com/uploads/pdf/magazine/2016/VB2016-Cobb.pdf

[24] ISACA. (2017). *State of cyber security 2017: Current trends in workforce development.* Schaumburg, Illinois: ISACA. Retrieved from https://cybersecurity.isaca.org/state-of-cybersecurity

[25] Center for Strategic and International Studies. (2016). *Hacking the skills shortage: A Study of the international shortage in cybersecurity skills.* Santa Clara, CA: Intel Security. Retrieved from https://www.mcafee.com/us/resources/reports/rp-hacking-skills-shortage.pdf

[26] Cisco. (2015). *Mitigating the cybersecurity skills shortage: Top insights and actions from Cisco Security Advisory Services.* Cisco. Retrieved from https://www.cisco.com/c/dam/en/us/products/collateral/security/cybersecurity-talent.pdf

[27] Hull, J. L. (2017). Analyst burnout in the cyber security operation center - CSOC: A phenomenological study (Doctoral dissertation). Retrieved from Proquest Dissertations and Theses database. (UMI No. 10282755)

CONVERSATIONAL TOPIC: ARTIFICIAL INTELLIGENCE

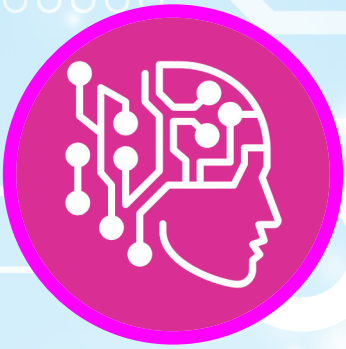

LogicHub

How to Measure and Improve Decision Automation for Cybersecurity (Next Gen SOAR)

By Kumar Saurabh

Facing an increasingly sophisticated barrage of threats, Security Operations Centers (SOCs) today are evaluating a variety of security tools, including security automation tools.

Security Operations, Automation and Response (SOAR) systems promise to automate data collection and threat remediation. They collect alerts and log data and, once an analyst has studied this data and decided upon a course of action, they perform some automated steps, such as closing ports and deleting files, to contain the threat. Some security automation platforms go even further, promising to completely automate the analysis and decision-making.

If these platforms could automate not just the comparatively simple steps of data collection and task automation, but also the more advanced work of threat analysis and decision-making that constitute the most difficult part of a security analyst's job, then security automation would bring an unprecedented level of efficiency and consistency to Security Operations. And analysts, who in nearly every SOC struggle to keep up with an unending torrent of alerts and user requests, would finally have more free time for proactive threat hunting and other critical – but often deferred – tasks.

Most SOC teams are hesitant and often times skeptical to automate more cognitive steps - steps that require domain knowledge, knowledge of tricks and techniques that they have acquired over the years, and expertise to know what data means and how to transform that data into a decision.

How do we objectively determine if an automation is capable of automating analysis and decision making?

ADDRESSING SOC SKEPTICS: CAN AUTOMATED DECISION ANALYSIS REACH THE RIGHT CONCLUSIONS?

An intelligent security automation platform can automate the cognitive work of analysts for any of a hundred or thousand types of threats. If the platform can do this job well, there's no reason for SOCs not to adopt this type of automation.

Hence, once we have automated a playbook, we want to measure and ensure that the automation is working as accurately as a human analyst would.

Undoubtedly, there will be some differences between what an automation will do from what an analyst would do. After all, sometimes two analysts in the same SOC will offer different conclusions and recommendations even when evaluating the same threat and working from the same playbook.

The goal isn't (and, practically, it can't be) to have automation match the conclusions of an analyst 100% of the time, since 1) even analysts differ in their conclusions, and 2) it's unlikely that machines and people will achieve the same results 100% of the time.

What SOCs really want to know is that any gap in accuracy between a human analyst and an automated system is within an acceptable threshold. The security platform and the analyst can differ – just not too much. In the event that they do differ by too much, further analysis is warranted before jumping to conclusions about the analyst, automation or both.

It's worth noting that even when automation cannot make a definitive conclusion itself about a threat, it can still automate a lot of an analyst's work, substantially reducing a SOC's overall workload.

With these preliminaries decided, let's consider how a SOC could go about methodically comparing the results delivered by a security automation platform to the results delivered by a senior analysts whose work in the SOC is considered exemplary.

What SOCs really want to know is that any gap in accuracy between a human analyst and an automated system is within an acceptable threshold.

METHODOLOGY: MEASURING RESULTS IN PHISHING TRIAGE

To compare the results of security automation and human analysis, it's helpful to focus on a single area or domain of security analysis.

Phishing triage is a common and – unfortunately – necessary part of daily life in just about any SOC. The goal of phishing triage is to determine which suspicious emails flagged by users or security tools are phishing attempts and which are benign.

Security automation can evaluate suspicious emails flagged as possible phishing attempts. As it analyzes emails, it sorts them into these distinct categories:

- Malicious (phishing)
- Benign (not phishing)
- Needs Manual Review (the system cannot automatically decide whether the email is a genuine phishing attempt.)

To benchmark the performance for this comparison, we will ask a human security analyst to triage a batch of emails (in this example, 30 emails) and sort them into two folders, Malicious and Benign. (We don't need a third category for emails needing manual review, since the analyst is already manually reviewing every email submitted for analysis.)

METRICS FOR EVALUATING RESULTS

After the security automation platform and the human analyst have both examined and sorted the complete batch of suspicious emails, we can compare the results. To make this comparison, we'll use these two metrics:

- **Accuracy**
  When the email was examined more closely or permitted to be delivered, did its classification turn out to be accurate?
- **Coverage**
  For what percentage of emails did the automation platform fully automate analysis and make a decision (as opposed to marking the email "Needs Manual Review")? Note that coverage refers to complete automation, obviating the need for an analyst to review

Ideally, both these metrics should be high. Let's examine why. Consider two possibilities:

If accuracy is around 90% or better and coverage is 75% or better, most SOCs we talk to would consider the accuracy of the automation to be good enough.

- **High accuracy, low coverage**
  If an automation platform delivered high accuracy (for example, 95%) but could reach a determination only in a limited number of cases (for example, covering only 50% of emails being considered), the automated platform would end up serving primarily as a labor-saving device. For example, coverage rate of 50% could eliminate half the SOC's analytical workload.
- **Low accuracy, high coverage**
  Conversely, if the platform delivered high coverage but low accuracy, the platform would not be useful at all to SOCs. Evaluating the majority of emails but miscategorizing them is hardly the result that any SOC is seeking.

Of the two metrics, accuracy is ultimately more important, because if accuracy is high, then the SOC can rely on automation to eliminate manual review of whatever percentage of emails is being covered and hope that that percentage will rise over time.

Ideally, both metrics should be high. Then SOCs would have a proven accurate solution for evaluating the majority of suspicious emails.

SUCCESS CRITERIA

If accuracy is around 90% or better and coverage is 75% or better, most SOCs we talk to would consider the accuracy of the automation to be good enough. That's because 75% of the analysts' workload has been reduced. And with the workload most security analysts are carrying today, that's a welcome relief.

What about accuracy rates? Results will vary obviously from playbook to playbook and SOC to SOC, but generally we find that:

- Security analyst error rates range from 16% to 25%.
- Automation platform error rates range from 1% to 5%.

Keep in mind that most first generation SOAR systems have a coverage rate of 0%, because they never fully automate their analysis. Instead, they require an analyst to stop work and review whatever data the SOAR is delivering an alert about.

When we say that coverage is 75%, we mean that three times out of four, the security analysts in the SOC need to take no action whatsoever in order for the security alert or incident to be fully resolved. That rate applies regardless of whether the resolution turns out to be dismissing the alert as a false positive or diagnosing it correctly as a specific type of threat and taking action to remediate it.

What if for a particular playbook or situation, the automation platform's accuracy or coverage is not good enough?

At LogicHub, when we encounter this situation, we spend an hour each week with a human analyst to understand why they would make the decision differently than our automation platform did. We will update the playbook within two business days, and then take another set of measurements over the next two days, and review the new metrics on the fifth day.

DARK LIGHT.ai
Active Defense Expert System

Thoughts about A.I. in Cybersecurity

By Shawn Riley

Generally when people hear the term AI they instantly think of data science derived AI such as machine learning and deep learning. This type of AI is very much needed as the amount of security data keeps increasing. When I started out in security in the early 1990s, the security analysts would have to manually go through the logs and other data sources to look for patterns of interest. As the amount of data grew, we started seeing data science being applied more and more across different data sets. This enabled data science approaches like machine learning to find the probable patterns and produce information for the analyst saying "these are probably the patterns you are looking for". I say probably because approaches like machine learning use probabilistic reasoning where the results are just conjecture until validated by a human with the necessary knowledge to understand the data they are looking at. It's often said the amount of data more than doubles every 2 years and this adds weight to why we need data science approaches like machine learning to do the preliminary analysis of the data for defenders. This also means the information, knowledge, and wisdom layers on top of the data are increasing as well.

As cybersecurity organizations deploy more and more sensors, they are also deploying more and more data science derived AI solutions to do that preliminary analysis. For the past several years this has been causing security analysts to drown in the information being produced in the same way they used to drown in the security data before wide spread use of data science derived AI solutions. The human analysts need to process the information being produced by all those solutions to verify the individual preliminary analysis results produced by algorithms to sort out the false positives from true detections. The problem is there is now far too much information being produced from the underlying data, when combined with information being shared by other organizations about threats and vulnerabilities, for most human security teams to process and take action on. The Ponemon Institute did a survey a few years ago that determined the average company has 75 security solutions, 96% of the information being produced wasn't being addressed, 19% were deemed reliable, 4% were actually investigated. The cybersecurity problems can't be addressed by data science derived AI alone. We also need knowledge engineering derived AI that focuses on organizing information into knowledge and can mimic how human security analysts and investigators apply the knowledge and wisdom contained in the knowledge-base.

The strength of knowledge engineering derived AI like modern expert systems is being able to mimic how human security analysts apply their knowledge and wisdom to the information to make sense of the preliminary analysis results coming from data science derived AI solutions and validating the performance of the point product producing the preliminary analysis results information.

Another strength of knowledge engineering derived AI is semantic interoperability. Which is the ability to integrate the information across different silos, that is in different formats and serializations into a common format and to organize the siloed information into integrated knowledge using W3C standardized ontologies (knowledge models created from knowledge representation language standards). This means the knowledge engineering derived AI can organize the information coming from the different data silos with knowledge from different frameworks such as MITRE ATT&CK, the NIST Cybersecurity Framework, NIST Cyber Resiliency Engineering Framework, ODNI Cyber Threat Framework, etc so the information is organized and can be looked at through the different lenses of knowledge frameworks and human mental models. Both data science derived AI and knowledge engineering derived AI are required pieces in the DHS and NSA sponsored Integrated Adaptive Cyber Defense (IACD) community.

Knowledge engineering derived AI really focuses on organizing and automating how human knowledge is applied to solve complex knowledge-driven automation challenges.

Knowledge engineering derived AI really focuses on organizing and automating how human knowledge is applied to solve complex knowledge-driven automation challenges. Human memory is both implicit and explicit. An example of implicit memory is that used to ride a bike. Whereas, explicit memory is made up of both Semantic (facts, dates, numbers, words) knowledge and Episodic (experiences) knowledge. Modern AI expert systems focus on encoding the explicit memory of humans and can capture semantic memory in knowledge models called ontologies and episodic memory in knowledge-driven playbooks that support deductive logical reasoning across semantic knowledge graphs.

Knowledge engineering derived AI expert systems can reason over the facts in the information it is looking at and based on those facts, infer (deduce) new facts into the investigation from the knowledge models (ontologies). It can also infer knowledge contained in the knowledge-base by following the knowledge-driven playbook workflow based on the encoded explicit knowledge of the human analyst. This is ideal for automating complex knowledge-driven processes that require the explicit knowledge and experience of human cyber defense analysts and investigators to make those processes scalable with the increasing amounts of data and information.

Applying AI in cybersecurity starts with knowing which type you need to solve the different problems faced by the security organization. If it's a data problem, then you need data science derived AI. If you've already invested in applying data science derived AI, then you're probably drowning in the information produced by the various data science derived AI solutions and don't have the humans you need to process, verify, and validate all the preliminary analysis results. You need to start thinking about investing in knowledge engineering derived AI solutions if you've reached this level. These are very different types of AI that don't have a lot of overlap but are extremely complimentary when both are used in the security enterprise. Data Science derived AI, Security Orchestration, and Knowledge Engineering derived AI

are 3 foundational technologies that are needed to support holistic security automation and to keep the human security team from drowning in the data and information.

It's important to remember that data science derived AI (aka non-symbolic AI) and knowledge engineering derived AI (aka symbolic AI) don't compete with each other but have a synergistic relationship, solving their own sets of problems as part of a holistic approach to applying AI in cybersecurity.

The State of Machine Learning in Cybersecurity

By TK Keanini

No matter where you stand on Machine Learning (ML) and Artificial Intelligence (AI), there's still plenty to talk about when it comes to how we as an industry are currently making use of them. With that in mind, I'd like to share some thoughts on ways we need to view machine learning and artificial intelligence as well as how we need to shift the conversation around them.

MORE EFFECTIVE = LESS OBVIOUS

I'm still amazed by how machine learning has remained a hot topic. That's not to say it doesn't deserve to be an area of interest though. Rather, what I'm suggesting is that what we should be talking about are the outcomes and capabilities it delivers. Some of you may remember when XML was a big deal, and everyone could not stop talking about it. Fast forward to today and no one advertises that they use XML since that would just be obvious and users care more about the functionality it enables. Machine learning will follow along the same path. In time, it will become an essential aspect of the way we approach security and become simply another background process. Once that happens, we can instead focus on talking about the analytical outcomes it enables.

AN ENSEMBLE CAST FEATURING MACHINE LEARNING

Anyone who has built an effective security analytics pipeline knows that job one is to ensure that it is resilient to active evasion. Threat actors know as much or more than you do about the detection methods within the environments they wish to penetrate and persist. The job of security analytics is to find the most stealthy and evasive threat actor activity in the network and to do this, you cannot just rely on a single technique. In order for that detection to happen, you need a diverse set of techniques that all complement one another. While a threat actor will be able to evade one or two of them simultaneously, they don't stand a chance against hundreds of them! The true power of detection is in diversity!

To better illustrate this idea, let's take a moment to consider a modern bank vault. Vaults employ a diverse set of detection techniques like motion, thermal, laser arrays, and on some physical dimension, an alarm will be tripped, and the appropriate response will ensue. We do the same in the digital world where machine learning helps us model timing or volumetric aspects of behaviors that are statistically normal and we can signal on outliers. This can be done all the way down at the protocol level where models are deterministic or all the way up to the application or users' behaviors which can sometimes be less deterministic. If you are in the same camp I am, we have had years to

refine these analytical techniques and have published well over 50 papers on the topic in the past 12 years.

THE PRECISION AND SCALE OF ML

So, at this point, you might be asking yourself some important questions. Why can't we just keep using lists of bad things and lists of good things? Why do we need machine learning in security analytics? What unique value does it bring us? The first thing I want to say here is that we should not be religious about machine learning or AI. To us, they are just another set of tools in the larger analytics pipeline. In fact, the most helpful analytics comes from using a bit of everything.

The first thing I want to say here is that we should not be religious about machine learning or AI. To us, they are just another set of tools in the larger analytics pipeline.

If you hand me a list and say, "If you ever see these patterns, let me know about it immediately!" I'm good with that. I can do that all day long and at very high speeds. But what if we are looking for something that cannot be known prior to the list making act? What if what we are looking for cannot be seen but only inferred? The shadows of the objects but never the objects if you will. What if we are not really sure what something is or the role it plays in the larger system (i.e., categorization and classification)? All these questions are where machine learning has contributed a great deal to security analytics. Let's point to a few examples.

THE ESSENCE OF ENCRYPTED TRAFFIC ANALYTICS

Encryption has made what was observable in the network impossible to observe. You can argue with me on this, but mathematics is not on your side, so let's just accept the fact that deep packet inspection is a thing of the past. We need a new strategy and that strategy is the power of inference. At Cisco, like other companies developing TLS fingerprinting techniques, we leverage the fact that all encrypted sessions begin unencrypted and that the routers and switches can send us an "Observable Derivative." This metadata coming from the network is a mathematical shadow of the payloads we cannot inspect directly because of encryption. Machine learning helps us train on these observable derivatives so that if its shape and size overtime is the same as some malicious behavior, we can bring this to your attention all without having to deal with decryption.

WHY IS THIS PRINTER BROWSING NETFLIX?

Sometimes we are lucky enough to know the identity and role of a user, application, or device as it interacts with systems across the network. The reality is, most days we are far from 100% on this, so machine learning can help us cluster network activity to make an assertion like, "based on the behavior and interactions of this thing, we can call it a printer!" When you are dealing with thousands upon thousands of computers interacting with one another across your digital business, even if you had a list at some point

in time – it is likely not up to date. The value to this labeling is not just so that you have objects with the most accurate labels, but so you can infer suspicious behavior based on its trusted role. For example, if a network device is labeled a printer, it is expected to act like a printer – future behavior can be expected from this device. If one day it starts to browse Netflix or checks out some code from a repository, it should be brought to your attention. With machine learning, you can infer from behavior what something is or if you already know what something is, you can predict its "normal" behavior and flag any behavior "not normal."

PATTERN MATCHING VERSUS BEHAVIORAL ANALYTICS

We cannot forget that no matter how fancy we get with the data science, if a human in the end will need to act on this information, they ultimately need to understand it.

Lists are great! Hand me a high-fidelity list and I will hand you back high-fidelity alerts generated from that list. Hand me a noisy or low fidelity list and I will hand you back noise. Back in 1959, computer gaming and AI pioneer Arthur Samuel defined machine learning as a "field of study that gives computers the ability to learn without being explicitly programmed." In security analytics, we can use it for just that and we can have analytical processes that implicitly program a list for you given the activity observed (the telemetry it is presented). Machine learning helps us implicitly put together a list that could not have been known a priori. In security, we complement what we know with what we can infer through negation. A simple example would be "if these are my sanctioned DNS servers and activities, then what is this other thing here?!" Logically, instead of saying something is A (or a member of set A), we are saying not-A but that only is practical if we have already closed off the world to {A, B} – not-A is B if the set is closed. If, however we did not close off the world to a fixed set of members, not-A could be anything in the universe which is not helpful.

USEFUL INFO FOR YOUR DAY-TO-DAY TASKS

I had gone my entire career measuring humans as if they were machines, and not measuring humans as humans. We cannot forget that no matter how fancy we get with the data science, if a human in the end will need to act on this information, they ultimately need to understand it. I had gone my entire career thinking that the data science could explain the results and while this is academically accurate, it is not helpful to the person who needs to understand the analytical outcome. The sense-making of the data is square in the domain of human understanding and this is why the only question we want to ask is "Was this alert helpful? Yes or no?" At the end of the day, we want to make sure that the person behind the console understands why an alert was triggered and if that helped them.

CONCLUSION

We owe a big round of applause to **artificial intelligence** for birthing the child we know and love named **machine learning** and all that it has contributed **to security analytics** over the past year. We remain pragmatic in its application as we know that, just because it is the new kid on the block, we cannot turn our backs on simple or complex lists of rules, simple statistical analysis, and any other method that has got us to where we are today.

Lucky for us, machine learning has already shown signs of playing well with its peers as we continue to find ways to improve existing security processes through pairing them with ML. It can't solve every single problem on its own, but when it works together with the people and processes that have come before it, we get that much closer to a more secure future. And if machine learning is the child of AI, who then are its brothers and sisters that we have yet to explore in Security Analytics?

CONTRIBUTOR BIOS

Bret Bergman

COO Partners in Performance, America

Bret has 25+ years of experience in electronics, high technology and consulting including many years in cyber security. He was the original VP/GM of Internet Security System's Managed Security and Knowledge Services Business Unit. He also 10 years running a semiconductor equipment manufacturer with facilities in U.S. and Korea during which time his firm consulted companies on microchip level security.

Joe Brule

Chair OASIS OpenC2 Technical Committee

Joe Brule has been working with NSA since 1997 and has focused on the Information Assurance mission since 2003. Currently, Mr. Brule is a cyber-engineer in the Capabilities Directorate and is a co-chair for the OASIS OpenC2 Technical Committee. His previous experience includes mission assurance for satellite systems, COMSEC engineering for space systems, Executive Secretary for the National Space INFOSEC Steering Council, and contributor to System Threat Assessment Reports and Capstone Threat Assessments (satellite systems and the global information grid). Mr. Brule has also contributed to CNSS Policy Number 12 (Space IA Policy) and was the primary author of a CNSS Memorandum (TRANSEC for Space Systems).

Geoff Hancock

CISO Advanced Cybersecurity Group

Geoff Hancock is the Chief Cybersecurity Executive at Advanced Cybersecurity Group responsible for strategy, operations and tactical implementation of cyber programs. He is a cybersecurity expert with more than 25 years' experience, has a bachelor's degree in computer science, is a CISSP, and has several other cybersecurity and business certifications. Prior to starting Advanced Cybersecurity Group, Mr. Hancock served in a number of executive cybersecurity and business roles, including CISO, CSO, GM Cyber business-technical solutions, Vice President Security Operations and CTO at several fortune 200 companies. Prior to that Mr. Hancock was in Joint Special Operations Command as a Special Operations Soldier specializing in Intelligence, Operations and Cybersecurity and worked across DoD and the Intelligence community on national security and cyber security issues.

He has spent his cyber career in technical and business operations in Military, Intelligence Community, Civilian agencies and many critical infrastructure sectors. Mr. Hancock has advised/co-authored the NIST Cybersecurity Framework and is an advisor to the Center for Internet Security and is an active contributor to the CIS 20 Critical Security Controls. During his career he has built or run over 20 Security Operations Centers including the Pentagon.

Mr. Hancock is an Adjunct professor at George Washington University where he co-created and teaches the World Cyber MBA Program. He is a guest lecturer on several technical cyber disciplines at NDU, NIU and UCF. He has authored/co-authored white papers, articles research papers and national policy recommendations.

He serves on the adviser council for CISO Magazine and the National Technology Security Coalition.

He is also chairman of the Federal CISO Advisory Council and advises several cybersecurity startups.

TK Keanini

Product Line CTO for Analytics Cisco

TK brings nearly 25 years of network and security experience to Cisco. With a penchant for driving technical innovation, he is responsible for integrating security solutions with private and public cloud-based computing platforms. He previously served as CTO for Lancope where he was responsible for leading the company's evolution toward integrating security solutions with private and public cloud-based computing platforms. Prior to that, he served as CTO of nCircle, driving product innovation that defined the vulnerability management and configuration compliance market. Before joining nCircle, he served as Vice President of Network Services for Morgan Stanley Online, where he built and secured a highly available online trading system. Previously, Keanini was a systems engineer at Cisco, advising top financial institutions on the design and architecture of their data networking infrastructures. He is a Certified Information Systems Security Professional (CISSP).

Philippe Langlois

Data Breach Investigations Report (DBIR) Author at Verizon

Philippe Langlois is currently a Data Breach Investigations Report Author at Verizon. For the last 3+ years he was Technical Product Manager for the CIS Critical Security Controls. In this role he leads an international community of cyber security experts who develop best practices known as the *CIS Critical Security Controls for Effective Cyber Defense*, a set of actions proven to mitigate 85% of the most prevalent cyber threats. He manages the production, writing, and publication of a range of cyber security resources. Working in collaboration with users of the CIS Critical Security Controls, he ensures the quality and utility of the Critical Controls guidance plus the availability of tools, scripts, and other resources aiding users with implementation of the Controls.

Previously he served as a Program Manager at the Multi-State Information Sharing and Analysis Center (MS-ISAC), within the Center for Internet Security. He managed the Nationwide Cybersecurity Review, establishing unique expertise in State, Local, Tribal and Territorial cyber security practice and assessment; co-chaired the Metrics, and Business Continuity/Recovery/Cyber Exercise Work Groups, and planned MS-ISAC sponsored exercises. He holds a Masters of Infrastructure Protection and International Security, a BA in Criminology and certifications as a Global Industrial Cyber Security Professional (GICSP), GIAC Penetration Tester (GPEN) and GIAC Critical Security Controls Certification (GCCC).

Kathy Lee Simunich

Computer Scientist Argonne National Laboratory

Kathy Lee Simunich is a Computer Scientist in the Strategic Security Sciences (SSS) division at Argonne National Laboratory. She has an M.S. Degree in Computer Science and a B.S. Degree in Meteorology. She has over three decades of experience in developing cross-platform model integration and simulation systems across many domains such as military logistics, environmental modeling systems, endangered species habitat management, real-time chemical and biological monitoring and analysis systems, health care procedures at hospitals, used nuclear fuel transportation and storage logistics, as well as the Cyber Fed Model, a communication framework for sharing Cyber Threat Information across the DOE, U.S. Government, and the North American Electrical Sector.

Aubrey Merchant-Dest

Federal CTO Symantec

Aubrey has 33 years of experience in Network & Cybersecurity Systems Engineering with both Carrier (fixed and mobile) and Enterprise environments. He came to Symantec through acquisition of Solera Networks which specialized in incident response and forensics where he was the Federal SE Manager. Prior SE positions included Qosmos/ ENEA, Ellacoya Networks, CloudShield, Springtide and iPolicy Networks with a focus on security, traffic engineering/

management and network analytics. He has an in-depth and hands-on understanding of networking from layer 2 through 7. With a total of 12 years in Deep Packet Inspection (DPI) and derivative technologies, his key focus is helping solve key issues related to network visibility/context, advancing workflow efficiency for cyber defenders through the fusion of network, application, user, and threat analytics. A key focus area is securing cloud adoption/migration. Aubrey is currently the Federal CTO at Symantec.

Kathleen M. Moriarty

Global Lead Security Architect EMC

Served as the IETF Security Area Director, Kathleen Moriarty is the Global Lead Security Architect with the EMC Office of the CTO (now Dell EMC) working on technology strategy and standards. Kathleen has been the primary author of multiple published standards and actively contributes to security standards activity in the IETF. Previously, as the Practice Manager for security consulting at EMC, Kathleen was responsible for oversight of key projects, and development of security programs, in addition to serving as the acting CISO of a global investment banking firm. Kathleen has also been the head of IT Security at MIT Lincoln Laboratory and the Director of Information Security at FactSet Research Systems. Kathleen holds a Master of Science degree in Computer Science from Rensselaer Polytechnic Institute.

Philip Reitinger

President and CEO, Global Cyber Alliance

Philip Reitinger has served as the President and CEO of the Global Cyber Alliance since December 2015. GCA is a non-profit organization focused on eradicating cybersecurity risks – risk by risk. Formerly he filled senior cybersecurity roles at VisionSpear LLC, Sony and Microsoft. In 2009 Mr. Reitinger was appointed as the Deputy Under Secretary for the National Protection and Programs Directorate at DHS. He also served as the first Executive Director of the DoD's Cyber Crime Center, and as Deputy Chief of the Computer Crime and Intellectual Property Section at DOJ. Mr. Reitinger has been awarded the Secretary of Homeland Security's Distinguished Service Medal and the Attorney General's John Marshall Award.

Shawn Riley

Chief Visionary Officer and Technical Advisor to the CEO, DarkLight Inc.

A recognized thought leader in the defense and intelligence communities, Shawn Riley, Chief Visionary Officer and Technical Advisor to the CEO at DarkLight Inc., brings over 25 years of cyber security, all source cyber threat intelligence, and artificial intelligence experience with an unparalleled understanding of the pitfalls that overtake modern security teams.

Tony Sager

Senior Vice President and Chief Evangelist, CIS

Tony Sager is a Senior Vice President and Chief Evangelist for CIS (The Center for Internet Security). In this role, he leads the development of the CIS Controls, a worldwide consensus project to find and support technical best practices in cybersecurity. Tony also serves as the Director of the SANS Innovation Center, a subsidiary of The SANS Institute.

Tony retired from the National Security Agency (NSA) after 34 years as an Information Assurance professional. He started his career in the Communications Security (COMSEC) Intern Program, and worked as a mathematical cryptographer and a software vulnerability analyst. In 2001, Tony led the release of NSA security guidance to the public. He also expanded the NSA's role in the development of open standards for security.

Mr. Sager holds a B.A. in Mathematics from Western Maryland College and an M.S. in Computer Science from The Johns Hopkins University.

Kumar Saurabh

CEO LogicHub

Kumar has 15 years of experience in the enterprise security and log management space leading product development efforts at ArcSight and SumoLogic. He has a passion for helping organizations improve the efficacy of their security operations, and personally witnessed the limitations of existing solutions in helping SOC analysts detect threats buried deep within mountains of alerts and events. This frustration led him to co-found LogicHub™ to empower cyber analysts by building intelligence automation, not just analytics.

Most recently Kumar was Co-founder and Vice President of Engineering at Sumo Logic. Previously, he was the data architect at Mint.com which was acquired by Intuit. Kumar was also one of the early engineering leads for the analytics and solutions team at ArcSight, and saw the company grow from zero revenue to IPO. Kumar earned his M.S. in Computer Science from Columbia University and B.S. in Computer Science from IIT Kharagpur.

Charles Schmidt

Group Leader The MITRE Corporation

Charles Schmidt is a Group Lead at the MITRE corporation, where he has worked for over 18 years in the field of cybersecurity. He has spent most of that time supporting security automation research and developing cybersecurity standards. He holds a Bachelor's degree in both Mathematics and Computer Science from Carleton College and a Master's degree in Computer Science from the University of Utah.

Kim Watson

IACD Technical Director JHU/APL

Kimberly K. Watson is a member of the Senior Staff at the Johns Hopkins Applied Physics Laboratory and is a Technical Director for Integrated Adaptive Cyber Defense (IACD). Ms. Watson was a technical leader at the Department of Homeland Security (DHS) from 2013-2015. Prior to her DHS role, she worked at NSA for more than 25 years, most of which was spent performing vulnerability discovery and technology evaluation activities.

Donnie Wendt

Security Engineer Mastercard

Donnie Wendt is a security engineer for MasterCard responsible for the design, architecture, and configuration of the security controls and monitoring protecting the MasterCard networks. His professional background includes over 30 years in information technology in various roles, including software development, network design, call center design, and project management. He joined MasterCard in 2004 as a Web engineer before moving into information security. He is also an adjunct professor of cybersecurity at Utica College.

Donnie earned a Master's degree in Cybersecurity from Utica College and a Bachelor's degree in Business Administration from Webster University. He maintains a Certified Information Systems Security Professional (CISSP) certification. He is currently pursuing a Doctorate of Science in Computer Science with an emphasis in Information Assurance and Cybersecurity from Colorado Technical University. His research focuses on security automation and adaptive cyber defense within the financial services industry.

NOTES

NOTES

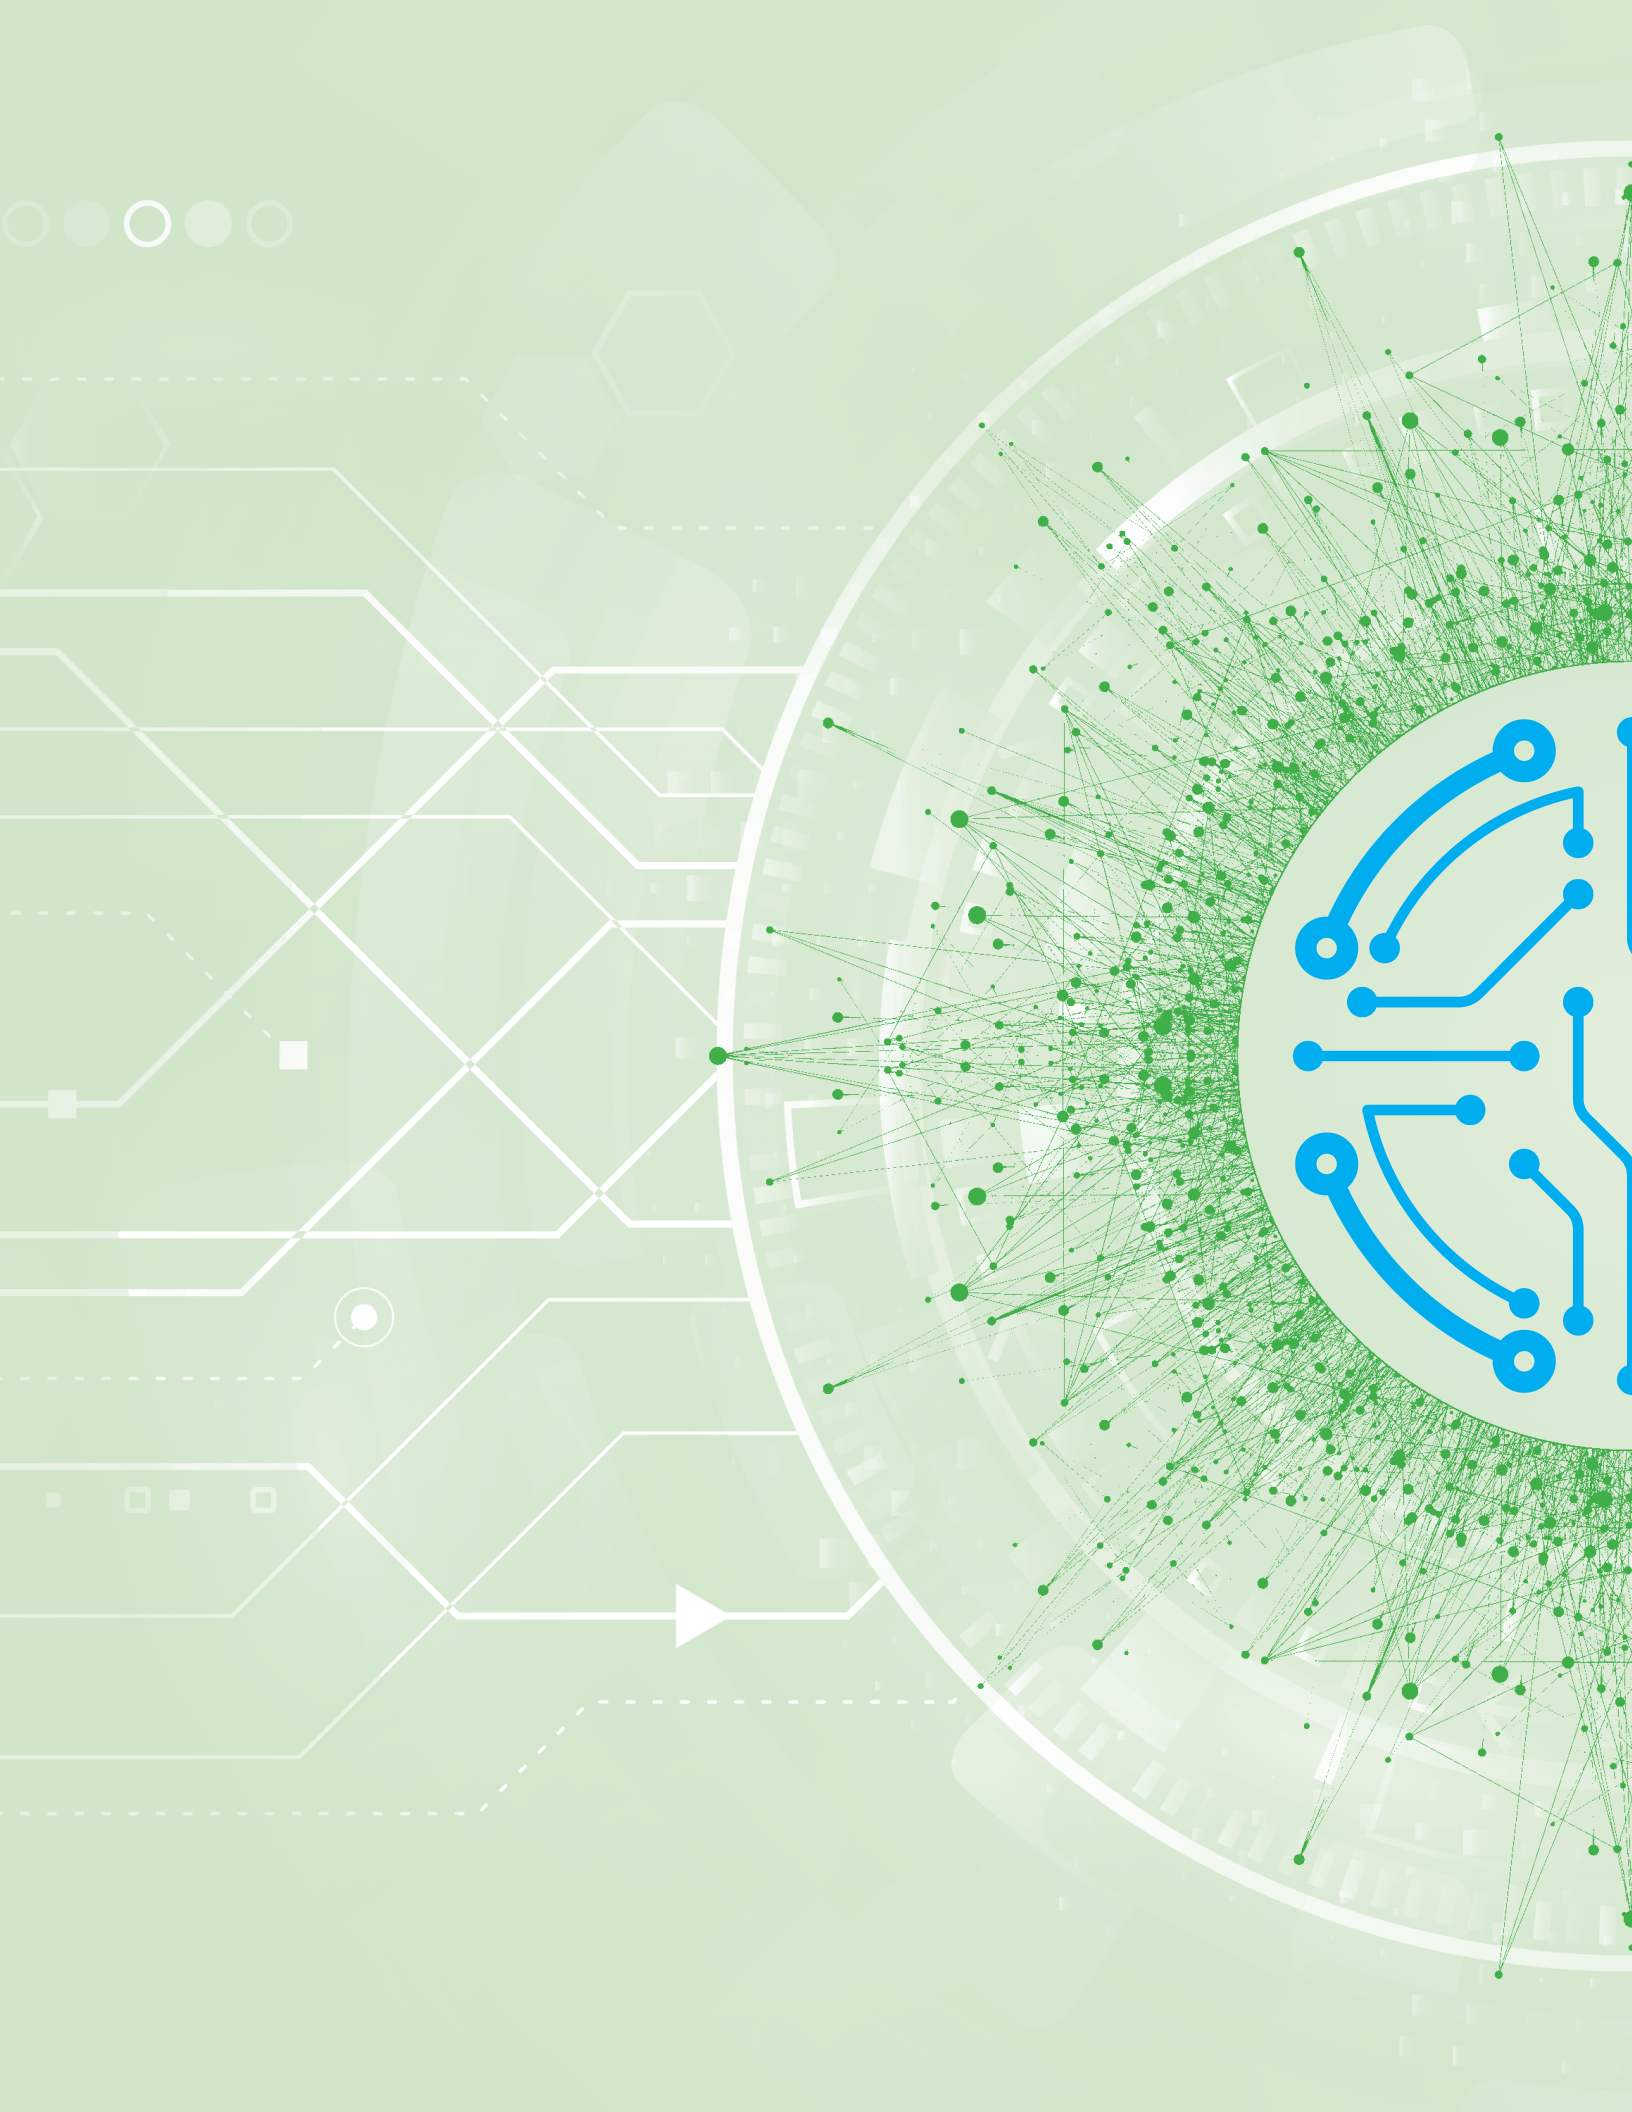

FCC_19-05816